Proven Infrastructure

EMC® VSPEX™ with Brocade Networking Solutions for END-USER COMPUTING

VMware Horizon View 5.3 and VMware vSphere for up to 2,000 Virtual Desktops

Enabled by Brocade VCS® Fabrics, EMC VNX, and EMC Next-Generation Backup

EMC VSPEX

Abstract

This document describes the EMC VSPEX validated with Brocade Networking Solutions for End-User Computing solution with VMware vSphere with EMC VNX for up to 2,000 virtual desktops.

May, 2014

BROCADE

EMC²

Copyright © 2014 EMC Corporation. All rights reserved. Published in the USA.

Published February 2014

EMC believes the information in this publication is accurate of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on the EMC online support website.

© 2014 Brocade Communications Systems, Inc. All Rights Reserved.

ADX, AnyIO, Brocade, Brocade Assurance, the B-wing symbol, DCX, Fabric OS, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, and Vyatta are registered trademarks, and HyperEdge, The Effortless Network, and The On-Demand Data Center are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

VMware Horizon View 5.3 and VMware vSphere for up to 2,000 Virtual Desktops Enabled by Brocade Network Fabrics, EMC VNX, and EMC Next-Generation Backup

Part Number:

Contents

Chapter 1 Executive Summary 15
Introduction 16
Audience 16
Document purpose 16
Business needs 17
Chapter 2 Solution Overview 18
Solution overview 19
Desktop broker 19
Virtualization 19
Compute 20
Network 20
Storage 21
EMC Next-Generation VNX 22
VNX performance 24
Virtualization management 26
Chapter 3 Solution Technology Overview 29
The technology solution 30
Key components 31
Desktop virtualization broker 32
Overview 32
VMware Horizon View 5.3 32
VMware View Composer 5.3 33
VMware View Persona Management 34
VMware View Storage Accelerator 34
Virtualization layer 35
VMware vSphere 35
VMware vCenter 35
VMware vSphere High Availability 35

EMC Virtual Storage Integrator for VMware 36
VNX VMware vStorage API for Array Integration Support 36
Compute layer 36
Network 38
File Storage Network with Brocade VDX Ethernet Fabric switches 38
Brocade VDX Ethernet Fabric Virtualization Automation Support 40
FC Block Storage Network with Brocade 6510 Fibre Channel switch 40
Storage 42
Overview 42
EMC VNX Series 42
EMC VNX Snapshots 43
EMC VNX SnapSure 43
EMC VNX Virtual Provisioning 44
VNX FAST Cache 49
VNX FAST VP (optional) 49
VNX file shares 49
ROBO 49
Backup and Recovery with EMC Avamar 50
Security layer 50
RSA SecurID Two-Factor Authentication 50
SecurID Authentication in the VSPEX End-User Computing for VMware Horizon View Environment 51
Required components 51
Compute, memory and storage resources 52
Other components 53
VMware vShield Endpoint 53
VMware vCenter operations manager for View 53
VMware Horizon File Share 54

Chapter 4 Solution Architectural Overview 59
Solution overview 60
Solution architecture 60
Logical architecture 60
Key components 62
VNX family storage arrays 65
Hardware resources 66
Software resources 69
Sizing for validated configuration 70

Server configuration guidelines 73
Overview 73
vSphere memory virtualization for VSPEX 74
Memory configuration guidelines 75
Brocade Network configuration guidelines 76
Overview 76
VLAN 78
Zoning (FC Block Storage Network only) 80
Storage configuration guidelines 81
Overview 81
vSphere Storage Virtualization for VSPEX 82
VSPEX storage building block 83
VSPEX end user computing validated maximums 85
High availability and failover 90
Introduction 90
Virtualization layer 90
Compute layer 91
Brocade Network layer 91
Storage layer 92
Validation test profile 94
Profile characteristics 94
Antivirus and antimalware platform profile 95
Platform characteristics 95
vShield Architecture 95
vCenter Operations Manager for View platform profile desktops 96
Platform characteristics 96
vCenter Operations Manager for View Architecture 97
Backup and recovery configuration guidelines 98
Sizing guidelines 98
Reference workload 98
Defining the reference workload 98
Applying the reference workload 99
Concurrency 99
Heavier desktop workloads 99
Implementing the reference architectures 100
Overview 100
Resource types 100

CPU resources 100
Memory resources 100
Network resources 101
Storage resources 102
Backup resources 102
Expanding existing VSPEX End User Ccomputing environments 102
Implementation summary 102
Quick assessment 103
Overview 103
CPU requirements 103
Memory requirements 103
Storage performance requirements 104
Storage capacity requirements 104
Determining equivalent reference virtual desktops 105
Fine tuning hardware resources 106
Chapter 5 VSPEX Configuration Guidelines 109
Configuration overview 110
Deployment process 110
Pre-deployment tasks 111
Overview 111
Deployment prerequisites 112
Customer configuration data 114
Prepare, connect, and configure Brocade network switches 114
Overview 114
Prepare Brocade Network Infrastructure 115
Configure Brocade VDX Ethernet Fabric Network (File based Storage Network) 115
Configure Brocade 6510 Fabric Network (Block based Storage Network) 116
Configure VLANs 118
Complete network cabling 118
Configure Brocade VDX 6740 Switch (File Storage) 119
Step 1: Verify and Apply Brocade VDX NOS Licenses 121
Step 2: Configure Logical Chassis VCS ID and RBridge ID 122
Step 3: Assign Switch Name 123
Step 4: Brocade VCS Fabric ISL Port Configuration 124
Step 5: Create vLAG for ESXi Hosts 127
Step 6: vCenter Integration for AMPP 129

Step 7: Create the vLAG for VNX ports....134
Step 8: Connecting the VCS Fabric to existing Infrastructure through Uplinks....136
Step 9 Configure MTU and Jumbo Frames (for NFS)....138
Step 10: Enable Flow Control Support....138
Step 11- Auto QOS for NAS....138
Prepare and configure the storage array....152
VNX configuration....152
Provision core data storage....153
Provision optional storage for user data....158
Provision optional storage for infrastructure virtual machines....161
Install and configure vSphere hosts....162
Overview....162
Install vSphere....162
Configure vSphere NetworkingConniknfnevjrevnervlkjnvrelgv....163
Jumbo frames....163
Connect VMware datastores....164
Plan virtual machine memory allocations....165
Install and configure SQL Server database....167
Overview....167
Create a virtual machine for Microsoft SQL Server....168
Install Microsoft Windows on the virtual machine....168
Install SQL Server....168
Configure the database for VMware vCenter....169
Configure database for VMware Update Manager....169
Configure database for VMware View Composer....169
Configure database for VMware Horizon View Manager....169
Configure the VMware Horizon View and View Composer database permissions....170
VMware vCenter Server Deployment....171
Overview....171
Create the vCenter host virtual machine....173
Install vCenter guest OS....173
Create vCenter ODBC connections....173
Install vCenter Server....173
Apply vSphere license keys....173
vStorage APIs for Array Integration (VAAI) Plug-in....173
Deploy PowerPath/VE (FC variant)....174

Install the EMC VSI plug-in 174
Set Up VMware View Connection Server 174
Overview 174
Install the VMware Horizon View Connection Server 176
Configure the View Event Log Database connection 176
Add a second View Connection Server 176
Configure the View Composer ODBC connection 176
Install View Composer 176
Link VMware Horizon View to vCenter and View Composer 176
Prepare master virtual machine 176
Configure View Persona Management Group Policies 177
Configure Folder Redirection Group Policies for Avamar 177
Configure View PCoIP Group Policies 178
Set Up EMC Avamar 178
Set up VMware vShield Endpoint 179
Overview 179
Verify desktop vShield Endpoint driver installation 180
Deploy vShield Manager appliance 180
Install the vSphere vShield Endpoint service 180
Deploy an antivirus solution management server 180
Deploy vSphere security virtual machines 180
Verify vShield Endpoint functionality 180
Set Up VMware vCenter Operations Manager for View 181
Overview 181
Create vSphere IP Pool for vCOps 182
Deploy vCenter Operations Manager vApp 182
Specify the vCenter server to monitor 182
Update virtual desktop settings 182
Create the virtual machine for the vCOps for View Adapter server 183
Install the vCOps for View Adapter software 183
Import the vCOps for View PAK File 183
Verify vCOps for View functionality 183
Summary 183

Chapter 6 Validating the Solution 185
Overview 186
Post-install checklist 187

Deploy and test a single virtual desktop 187
Verify the redundancy of the solution components 187
Provision remaining virtual desktops 188

Appendix A Bills of Materials 191
Bill of Materials for 500 virtual desktops 192
Bill of Material for 1,000 virtual desktops 194
Bill of Material for 2,000 virtual desktops 196

Appendix B Customer Configuration Data Sheet 199
Overview of customer configuration data sheets 200

Appendix C References 203
References 204
EMC documentation 204
Brocade Documentation 206
Other documentation 207

Appendix D About VSPEX 209
About VSPEX 210

Figures

Figure 1. Next-Generation VNX with multicore optimization 23
Figure 2. Active/active processors increase performance, resiliency, and efficiency 25
Figure 3. New Unisphere Management Suite 27
Figure 4. Solution components 30
Figure 5. Compute layer flexibility 37
Figure 6. Highly-available Brocade network design example – for File storage network 39
Figure 7. Example of Highly-Available Brocade network design – for FC block storage network 41
Figure 8. Storage pool rebalance progress 45
Figure 9. Thin LUN space utilization 46
Figure 10. Examining storage pool space utilization 47
Figure 11. Defining storage pool utilization thresholds 48
Figure 12. Defining automated notifications for block 48
Figure 13. Authentication control flow for View access requests originating on an external network 51
Figure 14. Logical architecture: VSPEX End-User Computing for VMware Horizon View with RSA 52
Figure 15. Horizon workspace architecture layout 55
Figure 16. Logical architecture: VSPEX End-User Computing for VMware View with Horizon Data 56
Figure 17. Logical architecture for block storage 61
Figure 18. Logical architecture for NFS storage 62
Figure 19. Hypervisor memory consumption 74
Figure 20. Required Brocade VDX network 79
Figure 21. Required networks with block storage variant 80
Figure 22. VMware virtual disk types 83
Figure 23. Storage layout building block for 500 virtual desktops 84
Figure 24. Storage layout building block for 1,000 virtual desktops 84

Figure 25. Core storage layout for 1,000 virtual desktops using VNX5400 .85
Figure 26. Optional storage layout for 1,000 virtual desktops using VNX5400 86
Figure 27. Core storage layout for 2,000 virtual desktops using VNX5600 .88
Figure 28. Optional storage layout for 2,000 virtual desktops using VNX5600 89
Figure 29. High availability at the virtualization layer 90
Figure 30. Redundant power supplies 91
Figure 31. Brocade Network layer High-Availability (VNX) – block storage network variant 92
Figure 32. Brocade Network layer High-Availability (VNX) - file storage network variant 92
Figure 33. VNX series high availability 93
Figure 34. Sample Ethernet network architecture 115
Figure 35. Sample network architecture – Block storage 117
Figure 36. Port types 125
Figure 37. Port Groups of the VDX 6740 126
Figure 38. Port Groups of the VDX 6740T and Brocade VDX 6740T-1G....126
Figure 39. VDX 6740 vLAG for ESXi hosts 127
Figure 40. VM Internal Network Properties 130
Figure 41. Example VCS/VDX network topology with Infrastructure connectivity 136
Figure 42. View all Data Mover parameters 155
Figure 43. Set nthread parameter 156
Figure 44. Storage System Properties dialog box 156
Figure 45. Create FAST Cache dialog box 157
Figure 46. Create Storage Pool dialog box Advanced tab 158
Figure 47. Storage Pool Properties dialog box Advanced tab 158
Figure 48. Storage Pool Properties dialog box 159
Figure 49. Manage Auto-Tiering Window 160
Figure 50. LUN Properties window 161
Figure 51. Virtual machine memory settings 166
Figure 52. View Composer Disks page 189

Tables

Table 1. VMX thresholds and settings 49
Table 2. Minimum hardware resources to support SecurID 52
Table 3. OVA virtual applications 54
Table 4. Minimum hardware resources to support VMware Horizon data 57
Table 5. Recommended EMC VNX storage needed for the Horizon Data NFS share 58
Table 6. Solution hardware 66
Table 7. Solution software 69
Table 8. Sample server configuration 71
Table 9. Server hardware 73
Table 10. Hardware resources for network infrastructure 76
Table 11. Storage hardware 81
Table 12. Number of disks required for different number of virtual desktops 85
Table 13. Validated environment profile 94
Table 14. Platform characteristics 95
Table 15. Platform characteristics 96
Table 16. Virtual desktop characteristics 99
Table 17. Blank worksheet row 103
Table 18. Reference virtual desktop resources 105
Table 19. Example worksheet row 105
Table 20. Example applications 106
Table 21. Server resource component totals 107
Table 22. Blank customer worksheet 108
Table 23. Deployment process overview 110
Table 24. Tasks for pre-deployment 111
Table 25. Deployment prerequisites checklist 112
Table 26. Tasks for switch and network configuration 114
Table 27. Brocade switch default settings 141

Table 28. Brocade 6510 FC switch Configuration Steps 141
Table 29. Brocade switch default settings 142
Table 30. Tasks for storage configuration 152
Table 31. Tasks for server installation 162
Table 32. Tasks for SQL Server database setup 167
Table 33. Tasks for vCenter configuration 171
Table 34. Tasks for VMware Horizon View Connection Server setup 175
Table 35. Tasks required to install and configure vShield Endpoint 179
Table 36. Tasks required to install and configure vCOps 181
Table 37. Tasks for testing the installation 186
Table 38. Common Server information 200
Table 39. vSphere Server information 201
Table 40. Array information 201
Table 41. Brocade Network infrastructure information 201
Table 42. VLAN information 202
Table 43. Service accounts 202

Chapter 1 Executive Summary

This chapter presents the following topics:

Introduction 16

Audience 16

Document purpose 16

Business needs 17

Introduction

VSPEX™ with Brocade networking solutions, validated and modular architectures are built with proven best-of-breed technologies to create complete virtualization solutions that enable you to make an informed decision in the hypervisor, compute, and networking layers. VSPEX eliminates server virtualization planning and configuration burdens. When embarking on server virtualization, virtual desktop deployment, or IT consolidation, VSPEX accelerates your IT Transformation by enabling faster deployments, choice, greater efficiency, and lower risk.

This document is a comprehensive guide to the technical aspects of this solution. Server capacity is provided in generic terms for required minimums of CPU, memory, and network interfaces; the customer is free to select the server hardware of their choice that meets or exceeds the stated minimums.

Audience

The readers of this document are expected to have the necessary training and background to install and configure an End-User Computing solution based on VMware View with VMware vSphere as a hypervisor, Brocade VDX Ethernet Fabric or Connectrix-B Fibre Channel series switches, EMC VNX series storage systems, and associated infrastructure as required by this implementation. External references are provided where applicable and the reader should be familiar with these documents.

Readers are also expected to be familiar with the infrastructure and database security policies of the customer installation.

Individuals focused on selling and sizing a VSPEX End-User Computing for VMware Horizon View solution should pay particular attention to the first four chapters of this document. After purchase, implementers of the solution should focus on the configuration guidelines in Chapter 5, the solution validation in Chapter 6, and the appropriate references and appendices.

Document purpose

This document is an initial introduction to the VSPEX End-User Computing architecture, an explanation of how to modify the architecture for specific engagements, and instructions on how to deploy the system.

The VSPEX with Brocade VDX End-User Computing architecture provides the customer with a modern system capable of hosting a large number of virtual desktops at a consistent performance level. This solution executes on the VMware vSphere virtualization layer backed by the highly available VNX storage family for storage and the VMware Horizon View desktop broker. The Compute and Network components are vendor definable,

redundant Brocade network switches and VNX storage family and sufficiently powerful to handle the processing and data needs of a large virtual machine environment.

The 500, 1,000, and 2,000 virtual desktop environments discussed are based on a defined desktop workload. While not every virtual desktop has the same requirements, this document contains methods and guidance to adjust your system to be cost effective when deployed. A smaller 250 virtual desktop environment based on the VNXe3300 is described in the document: *EMC VSPEX End-User Computing for VMware Horizon View 5.3 and VMware vSphere for up to 250 Virtual Desktops*.

An end-user computing or virtual desktop infrastructure is a complex system offering. This document facilitates its setup by providing up-front software and hardware material lists, systematic sizing guidance and worksheets, and verified deployment steps. After you install the last component, there are validation tests to ensure that your system is up and running properly. Following the instructions in this document will ensure an efficient and painless desktop deployment.

Business needs

Business applications are moving into a consolidated compute, network, and storage environment. EMC VSPEX with Brocade Networking End-User Computing solutions uses VMware to reduce the complexity of configuring the components of a traditional deployment model. The complexity of integration management is reduced while maintaining the application design and implementation options. Administration is unified, while process separation can be adequately controlled and monitored. The following are the business needs for the VSPEX End-User Computing for VMware architectures:

- Provide an end-to-end virtualization solution to utilize the capabilities of the unified infrastructure components.
- Provide a VSPEX End-User Computing for VMware Horizon View solution for efficiently virtualizing up to 2,000 virtual desktops for varied customer use cases.
- Provide a reliable, flexible, and scalable reference design.

Chapter 2 Solution Overview

This chapter presents the following topics:

Solution overview 19
Desktop broker 19
Virtualization 19
Compute 20
Network 20
Storage 21

Solution overview

The EMC VSPEX End-User Computing with Brocade networking solutions for VMware Horizon View on VMware vSphere provides a complete system architecture capable of supporting up to 2,000 virtual desktops with a redundant server/network topology and highly available storage. The core components that make up this particular solution are desktop broker, virtualization, compute, networking, and storage.

Desktop broker

View is the virtual desktop solution from VMware that allows virtual desktops to be run on the VMware vSphere virtualization environment. It allows for the centralization of desktop management and provides increased control for IT organizations. View allows end users to connect to their desktop from multiple devices across a network connection.

Virtualization

VMware vSphere is the leading virtualization platform in the industry. For years, it has provided flexibility and cost savings to end users by enabling the consolidation of large, inefficient server farms into nimble, reliable cloud infrastructures. The core VMware vSphere components are the VMware vSphere Hypervisor and the VMware vCenter Server for system management.

The VMware hypervisor runs on a dedicated server and allows multiple operating systems to execute on the system at one time as virtual desktops. These hypervisor systems can then be connected to operate in a clustered configuration. These clustered configurations are then managed as a larger resource pool through the vCenter product and allow for dynamic allocation of CPU, memory, network, and storage across the cluster.

Features like vMotion, which allows a virtual machine to move between different servers with no disruption to the operating system, and Distributed Resource Scheduler (DRS) which perform vMotions automatically to balance load, make vSphere a solid business choice.

With the release of vSphere 5.5, a VMware virtualized environment can host virtual machines with up to 64 virtual CPUs and 1 TB of virtual RAM.

Compute

VSPEX allows the flexibility of designing and implementing the vendor's choice of server components. The infrastructure must meet the following requirements:

- Sufficient RAM, cores and memory to support the required number and types of virtual machines
- Sufficient network connections to enable redundant connectivity to the system switches
- Excess capacity to withstand a server failure and failover in the environment

Network

Brocade VDX Ethernet Fabric and Fibre Channel Fabric switch technology enable the implementation of a high performance, efficient, and resilient networks validated with the VSPEX proven architectures. VSPEX with Brocade provides the storage network infrastructure the meets the requirements for redundant links for to existing infrastructure, compute, and storage with traffic isolation based on industry best practices. Brocade networking solutions provide an open standards based solution that unleashes the full potential of high-density server virtualization, private cloud architectures, and EMC VNX storage.

Brocade VDX Ethernet Fabrics networking solutions provides the following attributes:

- Offers flexibility to deploy 1000BASE-T and upgrade to 10GBASE-T for higher bandwidth
- Delivers high performance and reduces network congestion with 10 Gigabit Ethernet (GbE) ports, low latency, and 24 MB deep buffers
- Improves capacity with the ability to create up to a 160 GbE uplink with Brocade ISL Trunking
- Manages an entire multitenant Brocade VCS fabric as a single switch with Brocade VCS Logical Chassis
- Provides efficiently load-balanced multipathing at Layers 1, 2, and 3, including multiple Layer 3 gateways
- Simplifies Virtual Machine (VM) mobility and management with automated, dynamic port profile configuration and migration
- Supports Software-Defined Networking (SDN) technologies within data, control, and management planes

Brocade 6510 Fibre Channel Fabric is the purpose-built, data center-proven network infrastructure for storage, delivering unmatched reliability, simplicity, and 4/8/16 Gbps performance. The Brocade 6510 Fibre Channel Fabrics networking solutions provides the following attributes:

- Provides exceptional price/performance value, combining flexibility, simplicity, and enterprise-class functionality in a 48-port switch
- Enables fast, easy, and cost-effective scaling from 24 to 48 ports using Ports on Demand (PoD) capabilities
- Simplifies management through Brocade Fabric Vision technology, reducing operational costs and optimizing application performance
- Simplifies deployment and supports high-performance fabrics by using Brocade ClearLink Diagnostic Ports (D_Ports) to identify optic and cable issues
- Simplifies and accelerates deployment with the Brocade EZSwitchSetup wizard and Dynamic Fabric Provisioning (DFP)
- Maximizes availability with redundant, hot-pluggable components and non-disruptive software upgrades

Storage

The EMC VNX storage family is the number one shared storage platform in the industry. Its ability to provide both file and block access with a broad feature set make it an ideal choice for any End-User Computing implementation.

The VNX storage family components include the following, which are sized for the stated reference architecture workload:

- Host adapter ports (for block)—Provide host connectivity through fabric into the array
- Data Movers (for file)—Front-end appliances that provide file services to hosts (optional if providing CIFS/SMB, or NFS services)
- Storage processors (SPs)—The compute component of the storage array. SPs are used for all aspects of data moving into, out of, and between arrays.
- Disk drives—Disk spindles and solid state drives (SSDs) that contain the host/application data and their enclosures.

Note: The term Data Mover refers to a VNX hardware component, which has a CPU, memory, and input/output (I/O) ports. It enables the CIFS (SMB) and NFS protocols on the VNX array.

The desktop solutions described in this document are based on the EMC VNX5400™ and EMC VNX5600™ storage arrays respectively. The VNX5400 can support a maximum of 250 drives and the VNX5600 can host up to 500 drives.

The EMC VNX series supports a wide range of business class features ideal for the end-user computing cloud environment, including:

- EMC Fully Automated Storage Tiering for Virtual Pools (FAST VP™)
- EMC FAST™ Cache
- File-level data deduplication/compression
- Block deduplication
- Thin provisioning
- Replication
- Snapshots/checkpoints
- File-level retention
- Quota management

EMC Next-Generation VNX

The EMC VNX flash-optimized unified storage platform delivers innovation and enterprise capabilities for file, block, and object storage in a single, scalable, and easy-to-use solution. Ideal for mixed workloads in physical or virtual environments, VNX combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's virtualized application environments.

VNX includes many features and enhancements designed and built upon the first generation's success. These features and enhancements include:

- More capacity with multicore optimization with Multicore Cache, Multicore RAID, and Multicore FAST Cache (MCx™)
- Greater efficiency with a flash-optimized hybrid array
- Better protection by increasing application availability with active/active
- Easier administration and deployment by increasing productivity with new Unisphere® Management Suite

VSPEX is built with Next-Generation VNX to deliver even greater efficiency, performance, and scale than ever before.

Flash-optimized hybrid array

VNX is a flash-optimized hybrid array that provides automated tiering to deliver the best performance to your critical data, while intelligently moving less frequently accessed data to lower-cost disks.

In this hybrid approach, a small percentage of flash drives in the overall system can provide a high percentage of the overall IOPS. Flash-optimized VNX takes full advantage of the low latency of flash to deliver cost-saving

optimization and high performance scalability. The EMC Fully Automated Storage Tiering Suite (FAST Cache and FAST VP) tiers both block and file data across heterogeneous drives and boosts the most active data to the cache, ensuring that customers never have to make concessions in cost or performance.

FAST VP dynamically absorbs unpredicted spikes in system workloads. As that data ages and becomes less active over time, FAST VP tiers the data from high-performance to high-capacity drives automatically, based on customer-defined policies. This functionality has been enhanced with four times better detail and with new FAST VP solid-state disks (SSDs) based on enterprise multi-level cell (eMLC) technology to lower the cost per gigabyte. All VSPEX use cases benefit from the increased efficiency.

VSPEX Proven Infrastructures deliver private cloud, end-user computing, and virtualized application solutions. With VNX, customers can realize an even greater return on their investment. VNX provides out-of-band, block-based deduplication that can dramatically lower the costs of the flash tier.

VNX Intel MCx Code Path Optimization

The advent of flash technology has been a catalyst in making significant changes in the requirements of midrange storage systems. EMC redesigned the midrange storage platform to efficiently optimize multicore CPUs to provide the highest performing storage system at the lowest cost in the market.

MCx distributes all VNX data services across all cores—up to 32, as shown in Figure 1. The VNX series with MCx has dramatically improved the file performance for transactional applications like databases or virtual machines over network-attached storage (NAS).

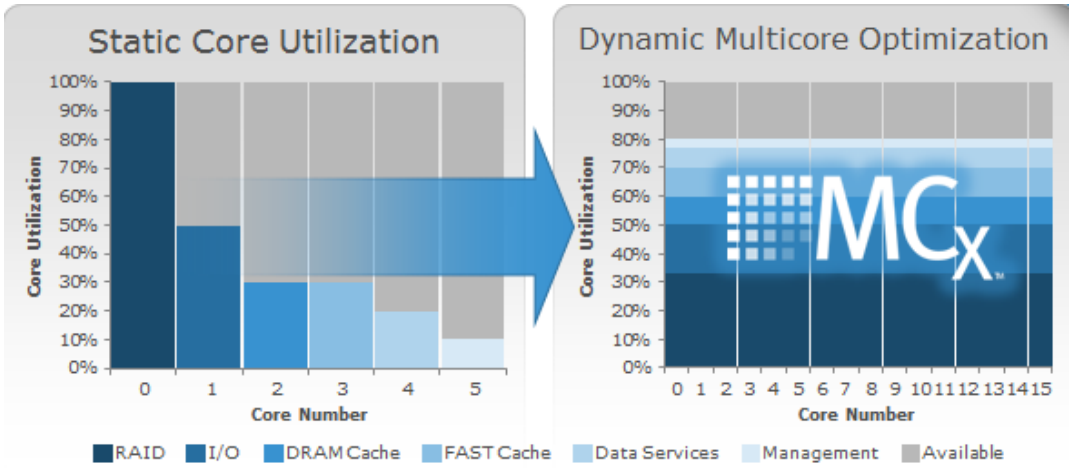

Figure 1. Next-Generation VNX with multicore optimization

Multicore Cache

The cache is the most valuable asset in the storage subsystem; its efficient use is the key to the overall efficiency of the platform in handling variable and changing workloads. The cache engine has been modularized to take advantage of all the cores available in the system.

Multicore RAID

Another important improvement to the MCx design is how it handles I/O to the permanent back-end storage—hard disk drives (HDDs) and SSDs. The modularization of the back-end data management processing, which enables MCx to seamlessly scale across all processors, greatly increases the performance of the VNX system.

VNX performance

VNX storage, enabled with the MCx architecture, is optimized for FLASH 1st and provides unprecedented overall performance; it optimizes transaction performance (cost per IOPS), bandwidth performance (cost per GB/s) with low latency, and capacity efficiency (cost per GB).

VNX provides the following performance improvements:

- Up to four times more file transactions when compared with dual controller arrays
- Increased file performance for transactional applications (for example, Microsoft Exchange on VMware over NFS) by up to three times with a 60 percent better response time
- Up to four times more Oracle and Microsoft SQL Server OLTP transactions
- Up to four times more virtual machines

Active/active array service processors

The new VNX architecture provides active/active array service processors, as shown in Figure 2, which eliminate application timeouts during path failover because both paths are actively serving I/O.

Load balancing is also improved, providing up to double the performance for applications. Active/active for block is ideal for applications that require the highest levels of availability and performance, but do not require tiering or efficiency services like compression, deduplication, or snapshot.

With this VNX release, VSPEX customers can use virtual Data Movers (VDMs) and VNX Replicator to perform automated and high-speed file-system migrations between systems. This process migrates all snaps and settings automatically, and enables the clients to continue operation during the migration.

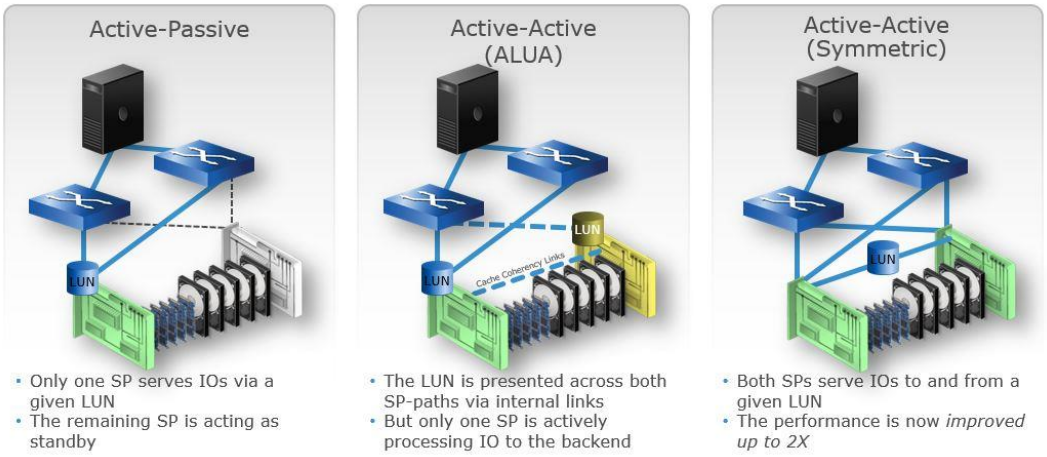
Active-Passive
LUN
• Only one SP serves IOs via a given LUN
• The remaining SP is acting as standby
Active-Active (ALUA)
LUN
LUN
Cache Coherency Links
• The LUN is presented across both SP-paths via internal links
• But only one SP is actively processing IO to the backend
Active-Active (Symmetric)
LUN
• Both SPs serve IOs to and from a given LUN
• The performance is now *improved up to 2X*

Figure 2. Active/active processors increase performance, resiliency, and efficiency

Virtualization management

VMware Virtual Storage Integrator

Virtual Storage Integrator (VSI) is a VMware vCenter plug-in that is available at no charge for VMware users with EMC storage. VSPEX customers can use VSI to simplify management of virtualized storage. VMware administrators can manage their VNX storage using the familiar vCenter interface.

With VSI, IT administrators can do more work in less time. VSI offers unmatched access control that enables you to efficiently manage and delegate storage tasks with confidence: you can perform daily management tasks with up to 90 percent fewer clicks and up to 10 times higher productivity.

VMware vStorage APIs for Array Integration

VMware vStorage application program interfaces (APIs) for Array Integration (VAAI) offload VMware storage-related functions from the server to the storage system, enabling more efficient use of server and network resources for increased performance and consolidation.

VMware vStorage APIs for Storage Awareness

VMware vStorage APIs for Storage Awareness (VASA) is a VMware-defined API that displays storage information through vCenter. Integration between VASA technology and VNX makes storage management in a virtualized environment a seamless experience.

Unisphere Management Suite

EMC Unisphere is the central management platform for the VNX series, providing a single, combined view of file and block systems, with all features and functions available through a common interface. Unisphere is optimized for virtual applications and provides industry-leading VMware integration, automatically discovering virtual machines and ESX servers and providing end-to-end, virtual-to-physical mapping. Unisphere also simplifies configuration of FAST Cache and FAST VP on VNX platforms.

The new Unisphere Management Suite extends Unisphere's easy-to-use, interface to include VNX Monitoring and Reporting for validating performance and anticipating capacity requirements. As shown in Figure 3, the suite also includes Unisphere Remote for centrally managing up to thousands of VNX and VNXe systems with new support for XtremSW Cache.

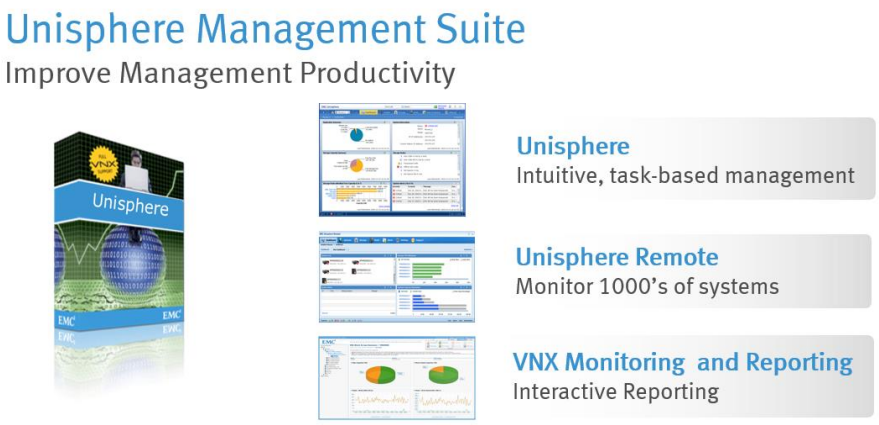

Figure 3. New Unisphere Management Suite

Chapter 3 Solution Technology Overview

This chapter presents the following topics:

The technology solution 30

Key components 31

Desktop virtualization broker 32

Virtualization layer 35

Compute layer 36

Network 38

Storage 42

Backup and Recovery with EMC Avamar 50

Security layer 50

Other components 53

The technology solution

This solution uses EMC VNX5400™ (for up to 1,000 virtual desktops) or VNX5600 (for up to 2,000 virtual desktops), Brocade Ethernet Fabric or Connectrix-B Fibre Channel switches, and VMware vSphere to provide the storage and computer resources for a VMware Horizon View environment of Windows 7 virtual desktops provisioned by VMware Horizon View™ Composer. Brocade Ethernet Fabric switches for File based storage network or Brocade 6510 Fibre Channel Fabric switches for Block storage network.

Figure 4 shows all of the computer resources and connections.

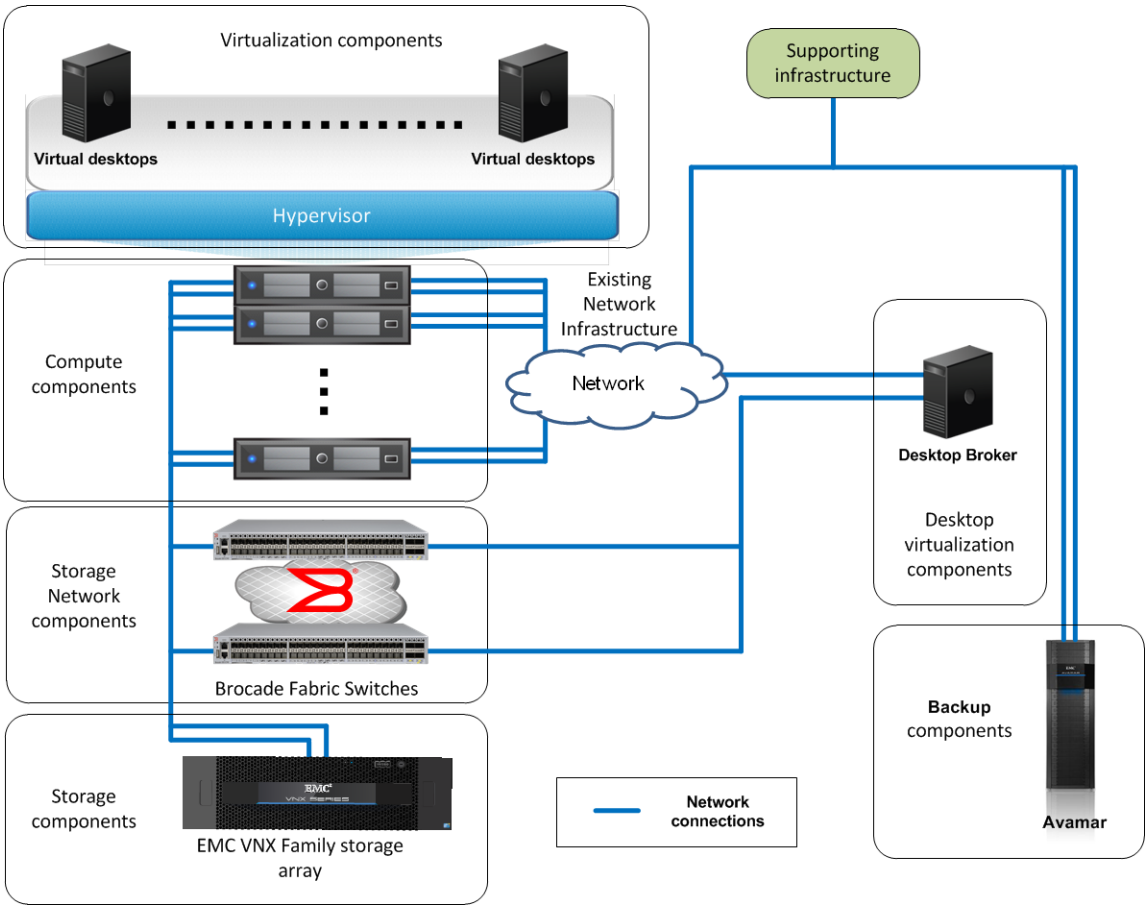

Figure 4. Solution components

Planning and designing the storage infrastructure for VMware Horizon View environment is a critical step because the shared storage must be able to absorb large bursts of I/O that occur over the course of a workday. These bursts can lead to periods of erratic and unpredictable virtual desktop performance. Users can adapt to slow performance, but unpredictable performance frustrates them and reduces efficiency.

To provide a predictable performance for an end-user computing solution, the storage system must be able to handle the peak I/O load from the clients while keeping response time to minimum. Designing for this workload involves the deployment of many disks to handle brief periods of extreme I/O pressure, which is expensive to implement. This solution uses EMC VNX FAST Cache to reduce the number of disks required.

EMC next-generation backup enables protection of user data and end-user recoverability. This is accomplished by using EMC Avamar® and its desktop client within the desktop image.

Key components

This section describes the key components of this solution.

- Desktop virtualization broker

  Manages the provisioning, allocation, maintenance, and eventual removal of the virtual desktop images that are provided to users of the system. This software is critical to enable on-demand creation of desktop images, allow maintenance to the image without affecting user productivity, and prevent the environment from growing in an unconstrained way.

- Virtualization layer

  Allows the physical implementation of resources to be decoupled from the applications that use them. In other words, the application's view of the resources available is no longer directly tied to the hardware. This enables many key features in the End-User Computing concept.

- Compute layer

Provides memory and processing resources for the virtualization layer software as well as for the applications running in the infrastructure. The VSPEX program defines the minimum amount of compute layer resources required but allows the customer to implement the requirements using any server hardware that meets these requirements.

- Network

Brocade VDX Ethernet Fabric or Connectrix-B Fibre Channel Fabric switches with Brocade Fabric networking technology connect the users of the environment to the resources they need and connects the storage layer to the compute layer. EMC VSPEX with Brocade networking solutions provides the required storage network connectivity, redundancy, and scalability. The EMC VSPEX with Brocade networking solutions enables the customer to implement a solution that provides a cost effective, resilient, and operationally efficient virtualization platform.

- Storage

  A critical resource for the implementation of the End-User Computing environment. Because of the way desktops are used, the storage layer must be able to absorb large bursts of activity as they occur without unduly affecting the user experience. This solution uses EMC VNX FAST Cache to efficiently handle this workload.

- Backup and Recovery

  Optional components of the solution that provide data protection in the event that the data in the primary system is deleted, damaged, or otherwise unusable.

- Security layerBackup and Recovery

  An optional solution component from RSA, which provides consumers with additional options to control access to the environment and ensure that only authorized users are permitted to use the system.

- Other Security layerBackup and Recovery

  Additional, optional, components that may improve the functionality of the solution exist, depending on the specifics of the environment.

- Solution architectureSecurity layerBackup and Recovery

  Provides details on all the components that make up the reference architecture.

Desktop virtualization broker

Overview

Desktop virtualization is a technology that encapsulates and delivers desktop services to a remote client device like thin clients, zero clients, smartphones, and tablets. It allows subscribers from different locations to access virtual desktops that are hosted on centralized computing resources at remote data centers.

In this solution, we used VMware Horizon View to provision, manage, broker, and monitor the desktop virtualization environment.

VMware Horizon View 5.3

VMware Horizon View 5.3 is a leading desktop virtualization solution that enables desktops to deliver cloud-computing services to users. VMware Horizon View 5.3 integrates effectively with vSphere 5.5 to provide:

- Performance optimization and tiered storage support—View Composer optimizes storage utilization and performance by reducing the footprint of virtual desktops. It also supports the use of different tiers of storage to maximize performance and reduce cost.

- Thin provisioning support—VMware Horizon View 5.3 enables efficient allocation of storage resources when virtual desktops are provisioned. This results in better utilization of the storage infrastructure and reduced capital expenditure (CAPEX)/operating expenditure (OPEX).
- Desktop virtual machine space reclamation—VMware Horizon View 5.3 can reclaim disk space that has been freed up within Windows 7 desktops. This feature ensures that the storage space required for the linked clone desktops is kept to a minimum throughout the lifecycle of the desktop.
- Requires the VMware Horizon View 5.3 Bundle. The VMware Horizon View Bundle includes access to all View features including vSphere Desktop, vCenter Server, View Manager, View Composer, View Persona Management, vShield Endpoint, VMware ThinApp, and VMware View Client with Local Mode.

The VMware Horizon View 5.3 release introduces the following enhancements for improving the user experience:

- A virtualized graphics processing unit (GPU) to support hardware-accelerated 3-D graphics
- Desktop access through HTML5 as well as the iOS and Android applications
- Support for Microsoft Windows 8

Note: What's New in VMware Horizon View 5.3 provides more details.

VMware View Composer 5.3

VMware View Composer 5.3 works directly with vCenter Server to deploy, customize, and maintain the state of the virtual desktops when using linked clones. Desktops provisioned as linked clones share a common base image within a desktop pool and therefore have a minimal storage footprint. The base image is shared among a large number of desktops. It is typically accessed with sufficient frequency to use EMC VNX FAST Cache, where frequently accessed data is promoted to flash drives. This behavior provides optimal I/O response time with fewer physical disks.

View Composer 5.3 also enables the following capabilities:

- Tiered storage support to enable the use of dedicated storage resources for the placement of both the read-only replica and linked-clone disk images
- An optional stand-alone View Composer server to minimize the impact of virtual desktop provisioning and maintenance operations on the vCenter server

This solution uses View Composer 5.3 to deploy dedicated virtual desktops running Windows 7 as linked clones.

VMware View Persona Management

VMware View Persona Management preserves user profiles and dynamically synchronizes them with a remote profile repository. View Persona Management does not require the configuration of Windows roaming profiles, eliminating the need to use Active Directory to manage View user profiles.

View Persona Management provides the following benefits over traditional Windows roaming profiles:

- With View Persona Management, View dynamically downloads a user's remote profile when the user logs in to a View desktop. View downloads persona information only when the user needs it.
- During login, View downloads only the files that Windows requires, such as user registry files. It then copies other files to the local desktop when the user or an application opens them from the local profile folder.
- View copies recent changes in the local profile to the remote repository at a configurable interval.
- During logout, View copies only the files that the user updated since the last replication to the remote repository.
- You can configure View Persona Management to store user profiles in a secure, centralized repository.

VMware View Storage Accelerator

View Storage Accelerator reduces the storage load associated with virtual desktops by caching the common blocks of desktop images into local vSphere host memory. Storage Accelerator uses a feature of the VMware vSphere platform called Content Based Read Cache (CBRC), which is implemented inside the vSphere hypervisor.

When enabled for the View virtual desktop pools, the host hypervisor scans the storage disk blocks to generate digests of the block contents. When these blocks are read into the hypervisor, they are cached in the host based CBRC. Subsequent reads of blocks with the same digest will be served from the in-memory cache directly. This significantly improves the performance of the virtual desktops, especially during boot storms, user login storms, or antivirus scanning storms, when a large number of blocks with identical content are read.

Virtualization layer

VMware vSphere

VMware vSphere is the market-leading virtualization platform that is used across thousands of IT environments around the world. VMware vSphere transforms a computer's physical resources by virtualizing the CPU, Memory, Storage, and Network. This transformation creates fully functional virtual desktops that run isolated and encapsulated operating systems and applications just like physical computers.

The high-availability features of VMware vSphere are coupled with DRS and VMware vMotion, which enables the seamless migration of virtual desktops from one vSphere server to another with minimal impact to the customer's usage.

This solution uses VMware vSphere Desktop Edition for deploying desktop virtualization. It provides the full range of features and functionalities of the vSphere Enterprise Plus edition, allowing customers to achieve scalability, high availability, and optimal performance for all of their desktop workloads. vSphere Desktop also comes with unlimited vRAM entitlement. vSphere Desktop edition is intended for customers who want to purchase only vSphere licenses to deploy desktop virtualization.

VMware vCenter

VMware vCenter is a centralized management platform for the VMware Virtual Infrastructure. It provides administrators with a single interface for all aspects of monitoring, managing, and maintaining the virtual infrastructure and can be accessed from multiple devices.

VMware vCenter is also responsible for managing some of the more advanced features of the VMware virtual infrastructure like VMware vSphere High Availability, DRS, vMotion, and Update Manager.

VMware vSphere High Availability

The VMware vSphere High Availability feature automatically allows the virtualization layer to restart virtual machines in various failure conditions.

- If the virtual machine operating system has an error, the virtual machine can be automatically restarted on the same hardware.
- If the physical hardware has an error, the impacted virtual machines can be automatically restarted on other servers in the cluster.

Note: In order to restart virtual machines on different hardware, those physical servers will need to have resources available. Follow the specific recommendations in Compute layer to enable this functionality.

VMware vSphere High Availability allows you to configure policies to determine which machines are restarted automatically and under what conditions these operations should be performed.

EMC Virtual Storage Integrator for VMware

EMC Virtual Storage Integrator (VSI) for VMware vSphere is a plug-in to the vSphere client that provides a single management interface that is used for managing EMC storage within the vSphere environment. Features can be added and removed from VSI independently, which provides flexibility for customizing VSI user environments. Use the VSI Feature Manager to manage the features. VSI provides a unified user experience that allows new features to be introduced rapidly in response to changing customer requirements.

The following features were used during the validation testing:

- Storage Viewer—extends the functionality of the vSphere Client to facilitate the discovery and identification of EMC VNX storage devices that are allocated to VMware vSphere hosts and virtual machines. Storage Viewer presents the underlying storage details to the virtual datacenter administrator, merging the data of several different storage mapping tools into a few seamless vSphere Client views.
- Unified Storage Management—Simplifies storage administration of the EMC VNX unified storage platform. It enables VMware administrators to seamlessly provision new Network File System (NFS) and Virtual Machine File System (VMFS) datastores and RDM volumes within the vSphere Client.

Refer to the EMC VSI for VMware vSphere product guides on EMC Online Support for more information.

VNX VMware vStorage API for Array Integration Support

Hardware acceleration with VMware vStorage API for Array Integration (VAAI) is a storage enhancement in vSphere that enables vSphere to offload specific storage operations to compatible storage hardware such as the VNX series platforms. With storage hardware assistance, vSphere performs these operations faster and consumes less CPU, memory, and storage fabric bandwidth.

Compute layer

The choice of a server platform for an EMC VSPEX infrastructure is not only based on the technical requirements of the environment, but also on the supportability of the platform, existing relationships with the server provider, advanced performance and management features, and many other factors. For this reason, EMC VSPEX solutions are designed to run on a wide variety of server platforms. Instead of requiring a given number of servers with a specific set of requirements, VSPEX documents a number of processor cores and an amount of RAM that must be achieved. This can be implemented with 2 servers (or 20) and still be considered the same VSPEX solution.

For example, let us assume that the compute layer requirements for a given implementation are 25 processor cores and 200 GB of RAM. One

customer might want to implement this by using white-box servers containing 16 processor cores, and 64 GB of RAM. A second customer chooses a higher-end server with 20 processor cores and 144 GB of RAM. Figure 5 depicts this example.

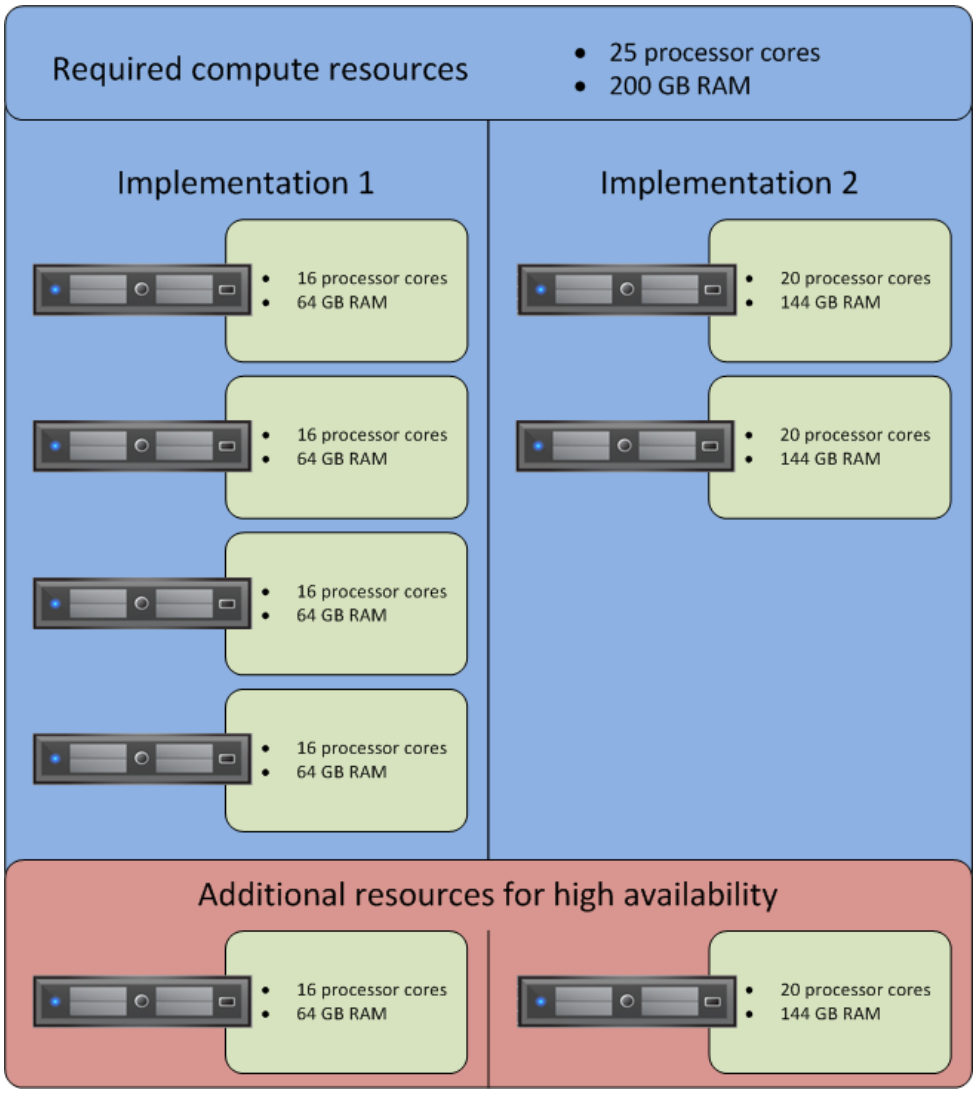

Figure 5. Compute layer flexibility

You should observe the following best practices in the compute layer:

- Use a number of identical or at least compatible servers. VSPEX implements hypervisor level high-availability technologies that may require similar instruction sets on the underlying physical hardware. By implementing VSPEX on identical server units, you can minimize compatibility problems in this area.
- If you are implementing hypervisor layer high availability, the largest virtual machine you can create will be constrained by the smallest physical server in the environment.

- Implement the high-availability features available in the virtualization layer to ensure that the compute layer has sufficient resources to accommodate at least single server failures. This allows you to implement minimal-downtime upgrades and tolerate single unit failures.

Within the boundaries of these recommendations and best practices, the compute layer for EMC VSPEX is flexible enough to meet your specific needs. The key constraint is that you provide sufficient processor cores and RAM per core to meet the needs of the target environment.

Network

VSPEX Proven Infrastructure with Brocade networking provides the dedicated storage network for host access to the VNX storage array. Brocade networking solutions provides options for block storage and file storage connectivity between compute and storage. The Brocade network is designed in the VSPEX reference architecture for block and file based storage traffic types to optimize throughput, manageability, application separation, high availability, and security. The storage network solution is implemented with redundant network links for each vSphere host, and VNX storage array. If a link is lost with any of the Brocade network infrastructure ports, the link fails over to another port. All network traffic is distributed across the active links.

The Brocade storage network infrastructure is deployed with redundant network links for each vSphere host, the storage array, the switch interconnect ports, and the switch uplink ports. This configuration provides both redundancy and additional storage network bandwidth. This configuration is also required regardless of whether the network infrastructure for the solution already exists or is being deployed alongside other components of the solution. Figure 13 provides an example of highly available network topology.

Note: The example is for IP-based networks, but the same underlying principles of multiple connections and eliminating single points of failure also apply to Fibre Channel based networks.

File Storage Network with Brocade VDX Ethernet Fabric switches

The Brocade VDX 6740 Ethernet Fabric series switches provide file based connectivity at 1 & 10 GbE in between the compute and VNX storage. Brocade® VDX with VCS Fabrics helps simplify networking infrastructures through innovative technologies and VSPEX file storage network topology design. The Brocade network validated solution uses virtual local area networks (VLANs) to segregate network traffic of VSPEX reference architecture for NFS storage traffic.

Brocade VDX 6740 switches support this strategy by simplifying network architecture while increasing network performance and resiliency with Ethernet fabrics. Brocade VDX with VCS Fabric technology supports active

– active links for all traffic from the virtualized compute servers to the EMC VNX storage arrays. The Brocade VDX provides a network with high availability and redundancy by using link aggregation for EMC VNX storage array. Figure 6 depicts an example of the Brocade network topology for file based storage.

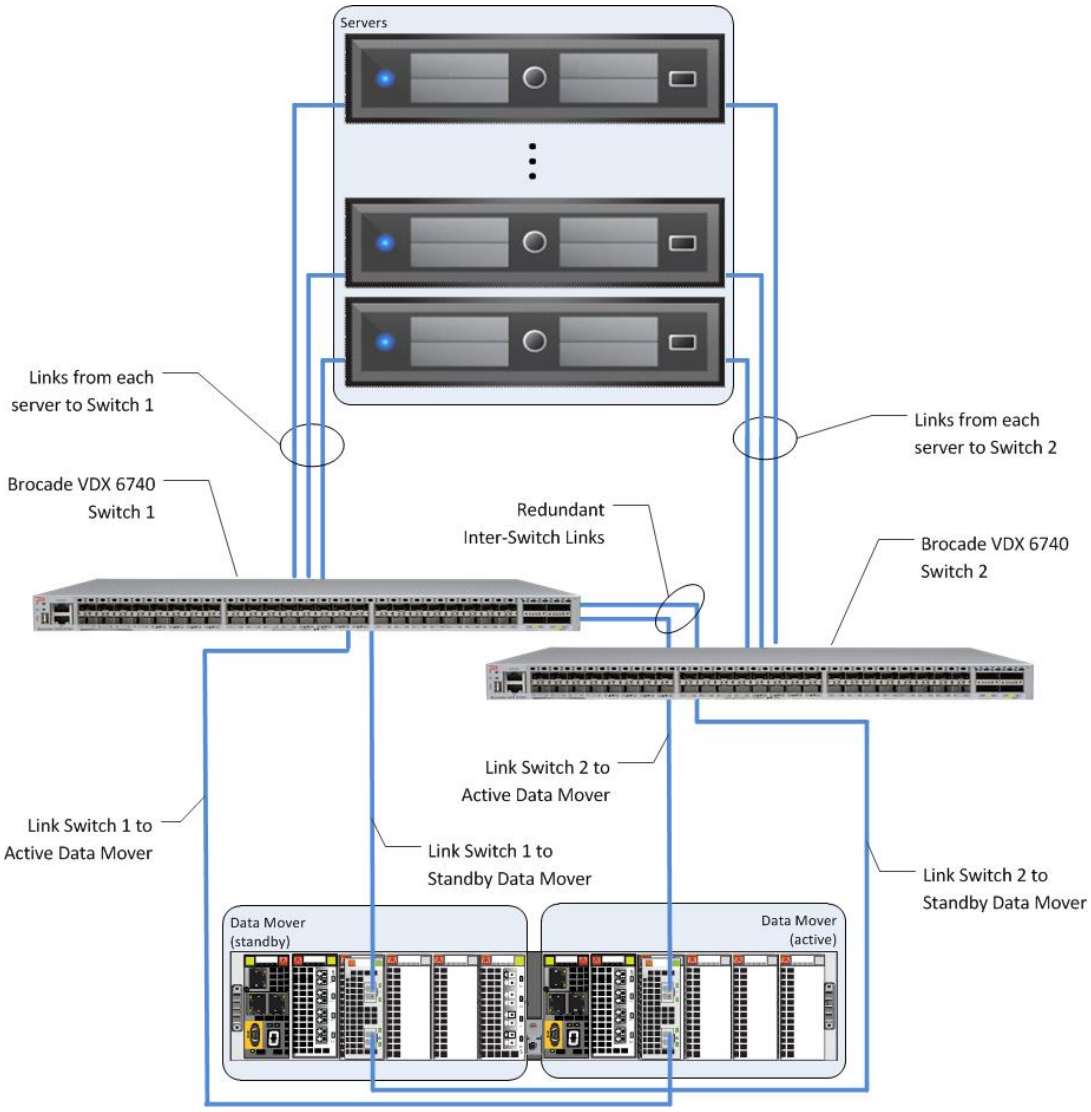
Servers
Links from each server to Switch 1
Links from each server to Switch 2
Brocade VDX 6740 Switch 1
Redundant Inter-Switch Links
Brocade VDX 6740 Switch 2
Link Switch 2 to Active Data Mover
Link Switch 1 to Active Data Mover
Link Switch 1 to Standby Data Mover
Link Switch 2 to Standby Data Mover
Data Mover (standby)
Data Mover (active)

Figure 6. Highly-available Brocade network design example – for File storage network

This validated solution uses virtual local area networks (VLANs) to segregate network traffic of various types to improve throughput, manageability, application separation, high availability, and security.

EMC unified storage platforms provide network high availability or redundancy by using link aggregation. Link aggregation enables multiple active Ethernet connections to appear as a single link with a single MAC address, and potentially multiple IP addresses. In this solution, Link Aggregation Control Protocol (LACP) is configured on VNX, combining multiple Ethernet ports into a single virtual device. If a link is lost in the

Ethernet port, the link fails over to another port. All network traffic is distributed across the active links.

Brocade VDX Ethernet Fabric Virtualization Automation Support

Brocade VDX with VCS Fabric technology offers unique features to support virtualized server and storage environments. Brocade VMaware network automation; for example, provides secure connectivity and full visibility to virtualized server resources with dynamic learning and activation of port profiles. By communicating directly with VMware vCenter, it eliminates manual configuration of port profiles and supports VMmobility across VCS fabrics within a data center.

FC Block Storage Network with Brocade 6510 Fibre Channel switch

The Brocade® 6510 with Gen 5 Fibre Channel Technology simplifies the storage network infrastructure through innovative technologies and supports the VSPEX highly virtualized topology design. The Brocade validated network solution simplifies server connectivity by deploying as full-fabric switch and enables fast, easy effective scaling from 24 to 48 Ports on Demand (PoD). The Brocade 6510 Fibre Channel switches maximizes availability with redundant architecture for Block Based storage traffic and hot-pluggable components and non-disruptive upgrades.

For block, the EMC VNX a unified storage platform is attached highly availability Brocade storage network redundancy by two ports per storage processor. If a link is lost on the storage processor front end port, the link fails over to another port. All storage network traffic is distributed across the active links.

Figure 7 depicts an example of the Brocade network topology for block based storage.

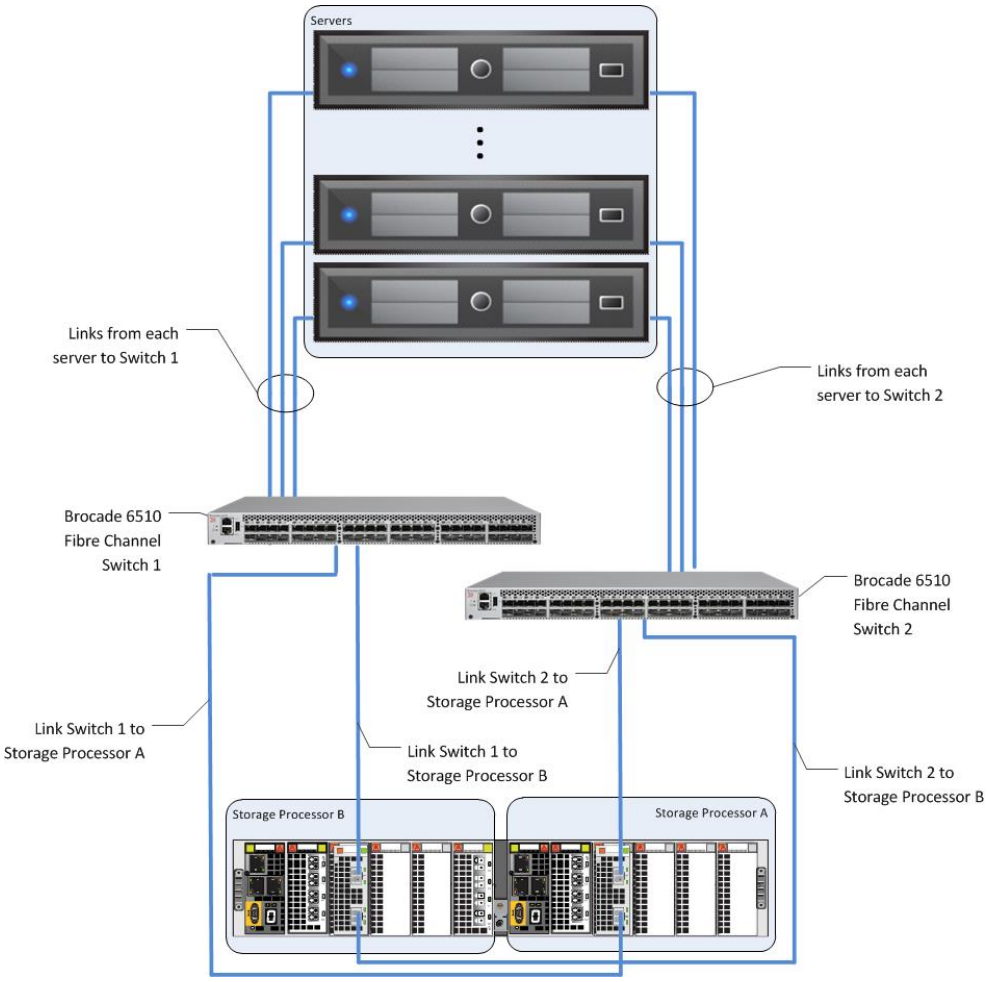
Servers
Links from each server to Switch 1
Links from each server to Switch 2
Brocade 6510 Fibre Channel Switch 1
Brocade 6510 Fibre Channel Switch 2
Link Switch 2 to Storage Processor A
Link Switch 1 to Storage Processor A
Link Switch 1 to Storage Processor B
Link Switch 2 to Storage Processor B
Storage Processor B
Storage Processor A

Figure 7. Example of Highly-Available Brocade network design – for FC block storage network

Brocade 6510 Fibre Channel switches supports provide high availabity for the VSPEX SAN infrastructure. Active – active links for all traffic from the virtualized to compute servers to the EMC VNX storage arrays. The Brocade® 6510 Switch meets the demands of hyper-scale, private cloud storage VSPEX storage traffic environments with market-leading Gen 5 Fibre Channel technology and capability that supports the VSPEX virtualized architecture.

The failure of a link in a route causes the network to reroute any traffic that was using that particular link—as long as an alternate path is available. Brocade Fabric Shortest Path First (FSPF) is a highly efficient routing algorithm that reroutes around failed links in less than a second.

ISL Trunking improves on this concept by helping to prevent the loss of the route. A link failure merely reduces the available bandwidth of the logical ISL trunk. In other words, a failure does not completely "break the pipe," but simply makes the pipe thinner. As a result, data traffic is much less likely to be affected by link failures, and the bandwidth automatically increases when the link is repaired.

Storage

Overview

The storage layer is a key component of any Cloud Infrastructure solution that serves data generated by applications and operating systems in a datacenter storage processing system. In this VSPEX solution, EMC VNX Series storage arrays are used to provide virtualization at the storage layer. This increases storage efficiency, management flexibility, and reduces total cost of ownership.

EMC VNX Series

The EMC Next-Generation VNX family is optimized for virtual applications and delivers innovation and enterprise capabilities for file and block storage in a scalable, easy-to-use solution. This storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's enterprises.

Intel Xeon processors power the VNX series for intelligent storage that automatically and efficiently scales in performance, while ensuring data integrity and security. VNX customer benefits include:

- Next-generation unified storage, optimized for virtualized applications
- Capacity optimization features including compression, deduplication, thin provisioning, and application-centric copies
- High availability, designed to deliver five 9s availability
- Automated Tiering with FAST VP and FAST Cache, which can be optimized for the highest system performance and lowest storage cost simultaneously
- Simplified management with EMC Unisphere for a single management interface for all NAS, SAN, and replication needs
- Up to three times better performance with the latest Intel Xeon multicore processor technology, optimized for flash

Various software suites and packs are also available for the VNX series, which provide multiple features for enhanced protection and performance

Software suites available

The following software suites are available:

- **FAST Suite**—Automatically optimizes for the highest system performance and the lowest storage cost simultaneously.
- **Local Protection Suite**—Practices safe data protection and repurposing.
- **Remote Protection Suite**—Protects data against localized failures, outages, and disasters.

- **Application Protection Suite**—Automates application copies and proves compliance.
- **Security and Compliance Suite**—Keeps data safe from changes, deletions, and malicious activity.

Software packs available

- **Total Efficiency Pack**—Includes all five software suites.
- **Total Protection Pack**—Includes local, remote, and application protection suites.

EMC VNX Snapshots

VNX Snapshots is a software feature that creates point-in-time data copies. You can use VNX Snapshots for data backups, software development and testing, repurposing, data validation, and local rapid restores. VNX Snapshots improves on the existing EMC VNX SnapView™ snapshot functionality by integrating with storage pools.

Note: LUNs created on physical RAID groups, also called RAID LUNs, support only SnapView snapshots. This limitation exists because VNX Snapshots requires pool space as part of its technology.

VNX Snapshots supports 256 writeable snapshots per pool LUN. It supports branching (also called 'snap of a snap'), as long as the total number of snapshots for any primary LUN is less than 256.

VNX Snapshots uses redirect on write (ROW) technology. ROW redirects new writes destined for the primary LUN to a new location in the storage pool. Such an implementation is different from copy on first write (COFW) used in SnapView, which holds the write activity to the primary LUN until the original data is copied to the reserved LUN pool to preserve a snapshot.

This release also supports consistency groups. You can combine several pool LUNs into a consistency group and snap them concurrently. When a snapshot of a consistency group is initiated, all writes to the member LUNs are held until the snapshots are created. Typically, consistency groups are used for LUNs that belong to the same application.

EMC VNX SnapSure

EMC VNX SnapSure™ is an EMC VNX Network Server software feature that enables you to create and manage checkpoints that are point in time, logical images of a production file system (PFS). SnapSure uses a copy-on-first-modify principle. A PFS consists of blocks. When a block within the PFS is modified, a copy containing the block's original contents is saved to a separate volume called the SavVol.

Subsequent changes made to the same block in the PFS are not copied into the SavVol. SnapSure reads the original blocks from the PFS in the SavVol and the unchanged PFS blocks remaining in the PFS according to a bitmap and block map data-tracking structure. These blocks combine to provide a complete point-in-time image called a checkpoint.

A checkpoint reflects the state of a PFS at the time the checkpoint is created. SnapSure supports the following checkpoint types:

- Read-only checkpoints—Read-only file systems created from a PFS
- Writeable checkpoints—Read/write file systems created from a read-only checkpoint

SnapSure can maintain a maximum of 96 read-only checkpoints and 16 writeable checkpoints per PFS, while allowing PFS applications continued access to real-time data.

Note: Each writeable checkpoint is associated with a read-only checkpoint, referred to as the *baseline* checkpoint. Each baseline checkpoint can have only one associated writeable checkpoint.

Using VNX SnapSure, available on EMC online support, provides more details.

EMC VNX Virtual Provisioning

EMC VNX Virtual Provisioning™ enables organizations to reduce storage costs by increasing capacity utilization, simplifying storage management, and reducing application downtime. Virtual Provisioning also helps companies to reduce power and cooling requirements and reduce capital expenditures.

Virtual Provisioning provides pool-based storage provisioning by implementing pool LUNs that can be either thin or thick. Thin LUNs provide on-demand storage that maximizes the utilization of your storage by allocating storage only as needed. Thick LUNs provide high performance and predictable performance for your applications. Both types of LUNs benefit from the ease-of-use features of pool-based provisioning.

Pools and pool LUNs are the building blocks for advanced data services such as FAST VP, advanced snapshots, and compression. Pool LUNs also support a variety of additional features, such as LUN shrink, online expansion, and user capacity threshold setting.

Virtual Provisioning allows you to expand the capacity of a storage pool from the Unisphere GUI after disks are physically attached to the system. VNX systems have the ability to rebalance allocated data elements across all member drives to use new drives after the pool is expanded. The rebalance function starts automatically and runs in the background after an "expand" action. You can monitor the progress of a rebalance operation from the **General** tab of the **Pool Properties** window in Unisphere, as shown in Figure 8.

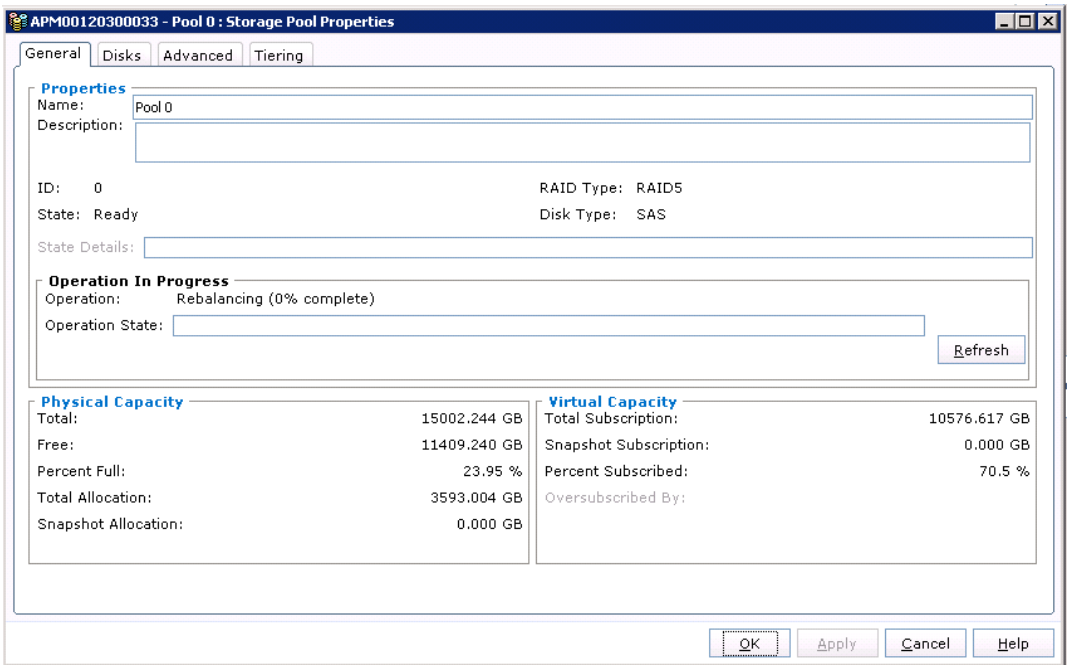

Figure 8. Storage pool rebalance progress

LUN expansion

Use pool LUN expansion to increase the capacity of existing LUNs. It allows for provisioning larger capacity as business needs grow.

The VNX family enables you to expand a pool LUN without disrupting user access. You can expand a pool LUN with a few simple clicks and the expanded capacity is immediately available. However, you cannot expand a pool LUN if it is part of a data-protection or LUN-migration operation. For example, you cannot expand snapshot LUNs or migrating LUNs.

LUN shrink

Use LUN shrink to reduce the capacity of existing thin LUNs. VNX can shrink a pool LUN. This capability is only available for LUNs served by Windows Server 2008 and later. The shrinking process involves these steps:

1. Shrink the file system from **Windows Disk Management**.
2. Shrink the pool LUN using a command window and the DISKRAID utility. The utility is available through the VDS provider, which is part of the EMC Solutions Enabler package.

The new LUN size appears as soon as the shrink process is complete. A background task reclaims the deleted or shrunk space and returns it to the storage pool. Once the task is complete, any other LUN in that pool can use the reclaimed space.

For detailed information on LUN expansion and shrinkage, refer to the *EMC VNX Virtual Provisioning - Applied Technology white paper,* available on EMC Online Support.

Alerting the user through the Capacity Threshold setting

You must configure proactive alerts when using a file system or storage pools based on thin pools. Monitor these resources so that storage is available for provisioning when needed to avoid capacity shortages.

Figure 8 explains why provisioning with thin pools requires monitoring.

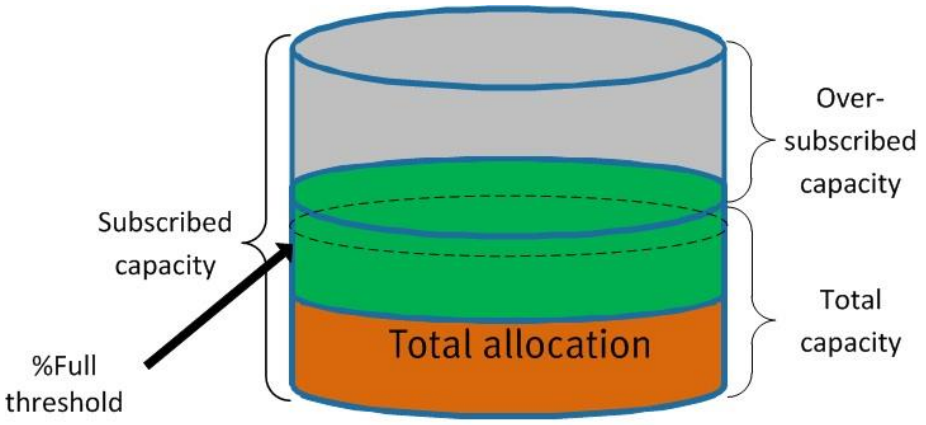

Figure 9. Thin LUN space utilization

Monitor the following values for thin pool utilization:

- **Total capacity**, which is the total physical capacity available to all LUNs in the pool
- **Total allocation**, which is the total physical capacity currently assigned to all pool LUNs
- **Subscribed capacity**, which is the total host-reported capacity supported by the pool
- Over**-subscribed capacity**, which is the amount of user capacity configured for LUNs that exceed the physical capacity in a pool

The total allocation must never exceed the total capacity, but if it approaches that point, add storage to the pools proactively before reaching a hard limit.

Figure 10 shows the **Storage Pool Properties** dialog box in Unisphere, which displays parameters such as Free, Percent Full, Total Allocation, Total Subscription of physical capacity, Percent Subscribed, and Oversubscribed for a system's virtual capacity.

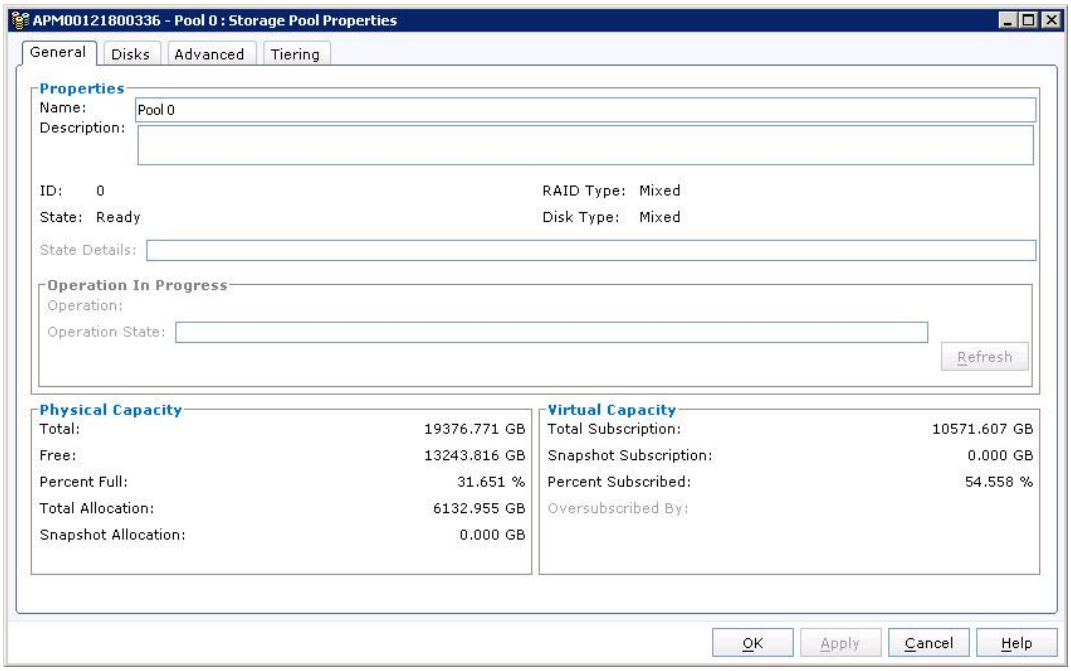

Figure 10. Examining storage pool space utilization

When storage pool capacity becomes exhausted, any requests for additional space allocation on thin-provisioned LUNs fail. Applications attempting to write data to these LUNs usually fail as well, and an outage is the likely result. To avoid this situation, you should:

1. Monitor pool utilization.
2. Set an alert that notifies you when thresholds are reached.
3. Set the **Percentage Full Threshold** to allow enough buffer space to correct the situation before an outage situation occurs.
4. Edit this setting by clicking **Advanced** in the **Storage Pool Properties** dialog box, as seen in Figure 11.

Note: This alert is only active if there are thin LUNs in the pool because thin LUNs provide the only way you can oversubscribe a pool. If the pool only contains thick LUNs, the alert is not active because there is no risk of running out of space due to oversubscription. You can also specify the value for **Percent Full Threshold**, which equals **Total Allocation/Total Capacity**, when a pool is created.

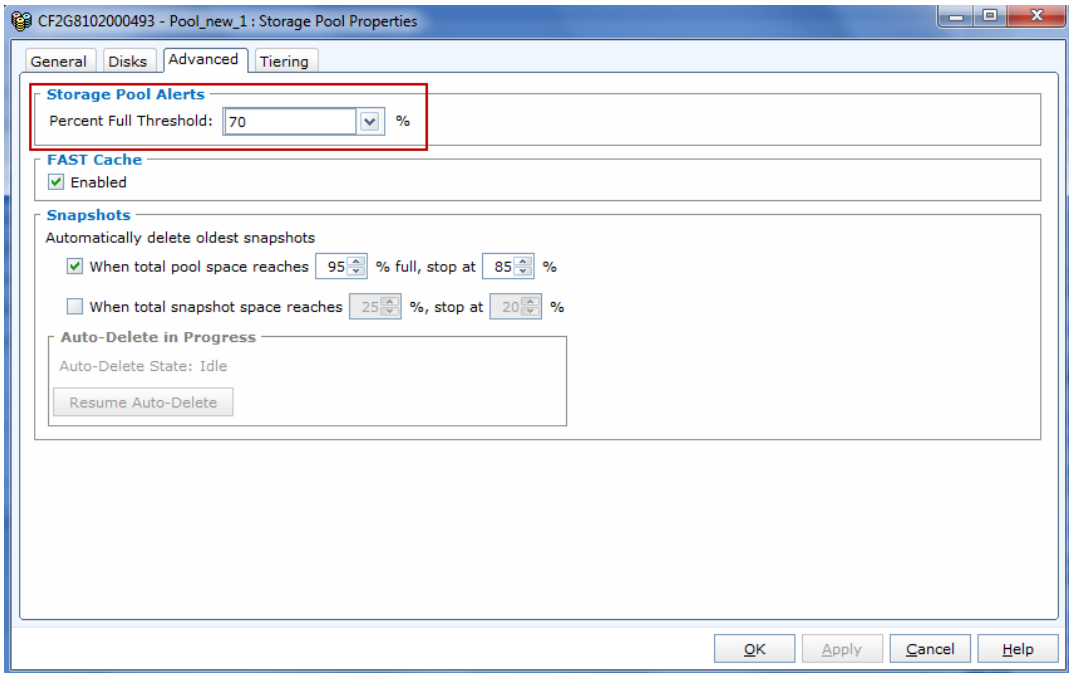

Figure 11. Defining storage pool utilization thresholds

Figure 12 shows the Unisphere Event Monitor Wizard, where you can view alerts. From this screen, you can also select the option to receive alerts through email, a paging service, or an SNMP trap.

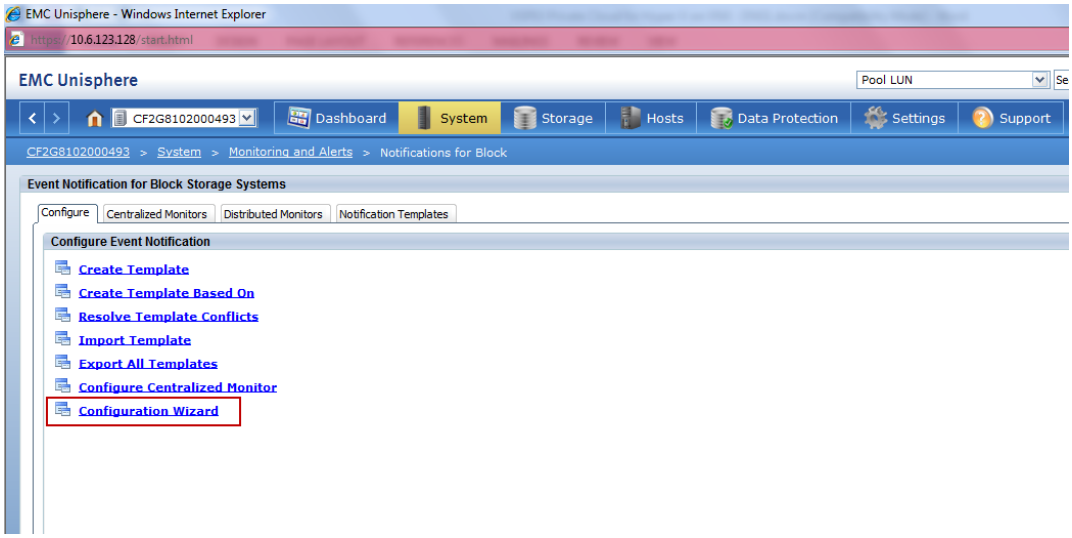

Figure 12. Defining automated notifications for block

Table 1 lists the information about thresholds and their settings for VNX Operating Environment (OE) for Block Release 33.

Table 1. VMX thresholds and settings

Threshold type	Threshold range	Threshold default	Alert severity	Side effect
User settable	1%-84%	70%	Warning	None
Built-in	N/A	85%	Critical	Clears user settable alert

Note: Allowing total allocation to exceed 90 percent of total capacity puts you at risk of running out of space and affecting all applications that use thin LUNs in the pool.

VNX FAST Cache

VNX FAST Cache, a part of the VNX FAST Suite, uses flash drives as an expanded cache layer for the array. FAST Cache is array-wide, non-disruptive cache available for both file and block storage. Frequently accessed data is copied to the FAST Cache in 64 KB increments. Subsequent reads and/or writes to the data chunk are serviced by FAST Cache. This enables immediate promotion of very active data to flash drives. This dramatically improves the response times for the active data and reduces data hot spots that can occur within the LUN.

VNX FAST VP (optional)

VNX FAST VP, a part of the VNX FAST Suite, enables you to automatically tier data across multiple types of drives to balance differences in performance and capacity. FAST VP is applied at the block storage pool level and automatically adjusts where data is stored based on how frequently it is accessed. Frequently accessed data is promoted to higher tiers of storage in 256 GB increments, while infrequently accessed data can be migrated to a lower tier for cost efficiency. This rebalancing of 256 GB data units, or slices, is done as part of a regularly scheduled maintenance operation.

VNX file shares

In many environments, it is important to have a common location in which to store files accessed by many users. CIFS or NFS file shares, available from a file server, provide this ability. VNX storage arrays can provide this service along with centralized management, client integration, advanced security options, and efficiency improvement features. For more information about VNX file shares, refer to *Configuring and Managing CIFS on VNX (300-013-429)* on EMC Online Support.

ROBO

Organizations with remote office and branch offices (ROBO) often prefer to locate data and applications close to the users in order to provide

better performance and lower latency. In these environments, IT departments need to balance the benefits of local support with the need to maintain central control. Local systems and storage should be easy for local personnel to administer, but also support remote management and flexible aggregation tools that minimize the demands on those local resources.

With VSPEX, you can accelerate the deployment of applications at remote offices and branch offices. Customers can also use **Unisphere Remote** to consolidate monitoring, system alerts, and reporting of hundreds of locations while maintaining operational simplicity and unified storage functionality for local managers.

Backup and Recovery with EMC Avamar

EMC Avamar delivers the protection confidence needed to accelerate deployment of VSPEX End User Computing.

Avamar empowers administrators to centrally backup and manage policies and EUC infrastructure components, while allowing end users to efficiently recover their own files from a simple and intuitive web-based interface. By moving only new, unique sub-file data segments, Avamar delivers fast daily full backups, with up to 90% reduction in backup times, while reducing the required daily network bandwidth by up to 99%. And all Avamar recoveries are single-step for simplicity.

Security layer

RSA SecurID Two-Factor Authentication

RSA SecurID two-factor authentication can provide enhanced security for the VSPEX End-User Computing environment by requiring the user to authenticate with two pieces of information, collectively called a passphrase, consisting of:

- Something the user knows: a PIN, which is used like any other PIN or password.
- Something the user has: A token code, provided by a physical or software "token," which changes every 60 seconds.

The typical use case deploys SecurID to authenticate users accessing protected resources from an external or public network. Access requests originating from within a secure network are authenticated by traditional mechanisms involving Active Directory or LDAP. A configuration description for implementing SecurID is available for the VSPEX End-User Computing infrastructures.

SecurID functionality is managed through RSA Authentication Manager, which also controls administrative functions such as token assignment to users, user management, and high availability.

SecurID Authentication in the VSPEX End-User Computing for VMware Horizon View Environment

SecurID support is built into VMware Horizon View, providing a simple activation process. Users accessing a SecurID-protected View environment will be initially authenticated with a SecurID passphrase, followed by normal authentication against Active Directory. In a typical deployment, one or more View Connection servers will be configured with SecurID for secure access from external or public networks, with other Connection servers accessed within the local network retaining Active Directory-only authentication.

Figure 13 depicts placement of the Authentication Manager server(s) in the View environment.

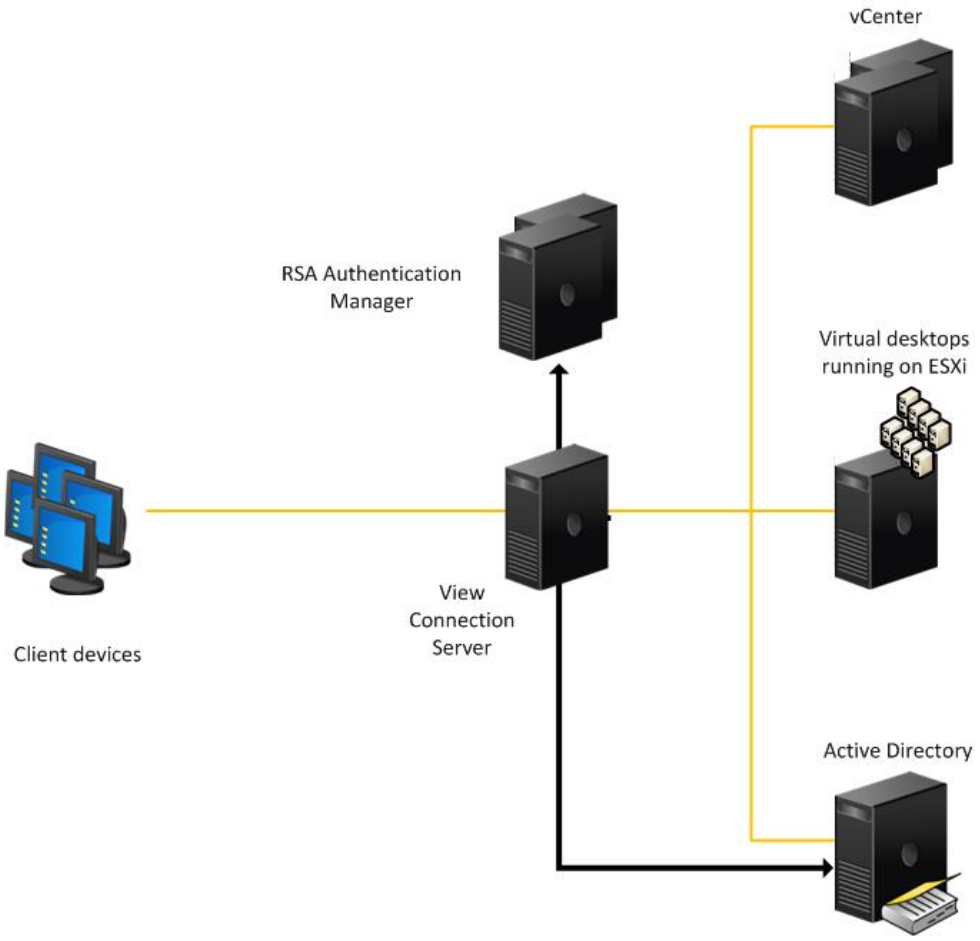

Figure 13. Authentication control flow for View access requests originating on an external network

Required components

Enablement of SecurID for VSPEX is described in Securing EMC VSPEX End-User Computing with RSA SecurID: VMware Horizon View 5.3 and VMware vSphere 5.5 for up to 2,000 Virtual Desktops Design Guide. The following components are required:

- **RSA SecurID Authentication Manager** (version 8.0)—Used to configure and manage the SecurID environment and assign tokens to users, Authentication Manager 8.0 is available as a virtual appliance running on VMware ESXi.
- **SecurID tokens for all users—**SecurID requires something the user knows (a PIN) with a constantly changing code from a "token" the user has in possession. SecurID tokens can be physical—displaying a new code every 60 seconds, which the user must then enter with a PIN—or software-based, wherein the user supplies a PIN and the token code is supplied programmatically. Hardware and software tokens are registered with Authentication Manager through "token records" supplied on a CD or other media.

Compute, memory and storage resources

Figure 14 depicts the VSPEX End-User Computing for VMware Horizon View environment with two infrastructure virtual machines added to support Authentication Manager. Table 2 shows the server resources needed; requirements are minimal and can be drawn from the overall infrastructure resource pool.

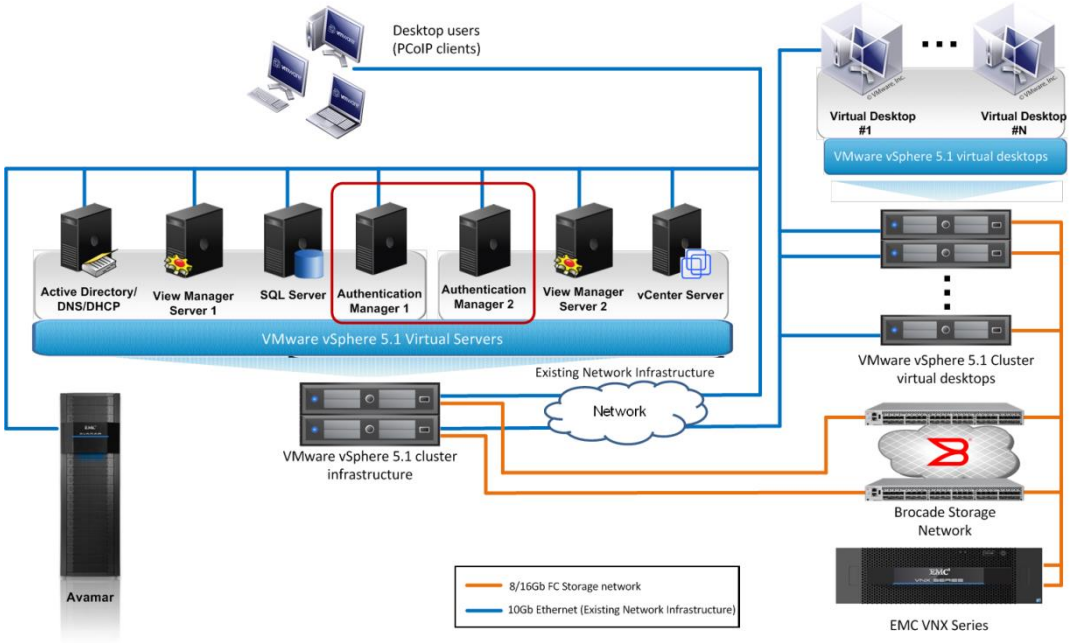

Figure 14. Logical architecture: VSPEX End-User Computing for VMware Horizon View with RSA

Table 2. Minimum hardware resources to support SecurID

	CPU (cores)	Memory (GB)	Disk (GB)	Reference
RSA Authentication Manager	2	4	100	RSA Authentication Manager 8.0 Performance and Scalability Guide

Other components

VMware vShield Endpoint

VMware vShield Endpoint offloads virtual desktop antivirus and antimalware scanning operations to a dedicated secure virtual appliance delivered by VMware partners. Offloading scanning operations improves desktop consolidation ratios and performance by eliminating antivirus storms; streamlining antivirus and antimalware deployment; and monitoring and satisfying compliance and audit requirements through detailed logging of antivirus and antimalware activities.

VMware vCenter operations manager for View

VMware vCenter Operations Manager for View provides end-to-end visibility into the health, performance, and efficiency of virtual desktop infrastructure (VDI). It enables desktop administrators to proactively ensure the best end-user experience, avert incidents, and eliminate bottlenecks. Designed for VMware Horizon View, this optimized version of vCenter Operations Manager improves IT productivity and lowers the cost of owning and operating VDI environments.

Traditional operations-management tools and processes are inadequate for managing large View deployments, because:

- The amount of monitoring data and quantity of alerts overwhelm desktop and infrastructure administrators.
- Traditional tools provide only a silo view and do not adapt to the behavior of specific environments.
- End users are often first to report incidents. They also troubleshoot performance problems that lead to fire drills among infrastructure teams, helpless help-desk administrators, and frustrated users.
- Lack of end-to-end visibility into the performance and health of the entire stack—including servers, storage, and networking—stalls large VDI deployments.
- IT productivity suffers from reactive management and the inability to proactively ensure quality of service.

VMware vCenter Operations Manager for View addresses these challenges and delivers higher team productivity, lower operating expenses, and improved infrastructure utilization.

Key features include:

- Patented self-learning analytics that adapt to your environment and continuously analyze thousands of metrics for server, storage, networking, and end-user performance
- Comprehensive dashboards that simplify monitoring of health and performance, identify bottlenecks, and improve infrastructure efficiency of your entire View environment

- Dynamic thresholds and smart alerts that notify administrators early in the process and provide more-specific information about impending performance issues
- Automated root-cause analysis, session lookup, and event correlation for faster troubleshooting of end-user problems
- Integrated approach to performance, capacity and configuration management that supports holistic management of VDI operations
- Design and optimizations specifically for VMware Horizon View
- Availability as a virtual appliance for faster time to value

VMware Horizon File Share

Horizon Data is a component of VMware Horizon Workspace that combines applications and data into a single, aggregated workspace. Horizon Data provides the flexibility for employees to access the data, no matter where that data resides, reducing the complexity of IT administration. This solution requires the presence of Active Directory (AD) and Domain Name Resolution (DNS).

Key Components

Horizon Workspace is distributed as a vApp, or an Open Virtual Appliance (.OVA) file, which can be deployed through VMware vCenter. Figure 15 shows the architecture of a basic Horizon Workspace layout. The OVA file contains the Virtual Appliances (VA) described in Table 3:

Table 3. OVA virtual applications

Application	Description
Configurator (configurator-va)	The Configurator appliance provides the central wizard UI and distributes settings across all other appliances in the vApp. It also provides the function of change the Network Gateway, vCenter, and SMTP settings.
Connector (connector-va)	The Connector appliance provides user authentication services; it can also bind with an Active Directory and synchronize under a defined schedule.
Manager (service-va)	The Manager appliance provides the web-based Horizon Workspace admin user interface, which controls the application catalog, user entitlements, workspace groups, and reporting service.
Data (data-va)	The Data appliance provides the service that allows you to store and share user files. It includes a web-based interface for previewing and performing functions on the user file.

Application	Description
Gateway (gateway-va)	The Gateway appliance enables a single, user-facing domain access to Horizon Workspace. As the central aggregation point for all user connections, the Gateway routes requests to the appropriate destination and proxies requests on behalf of user connections.

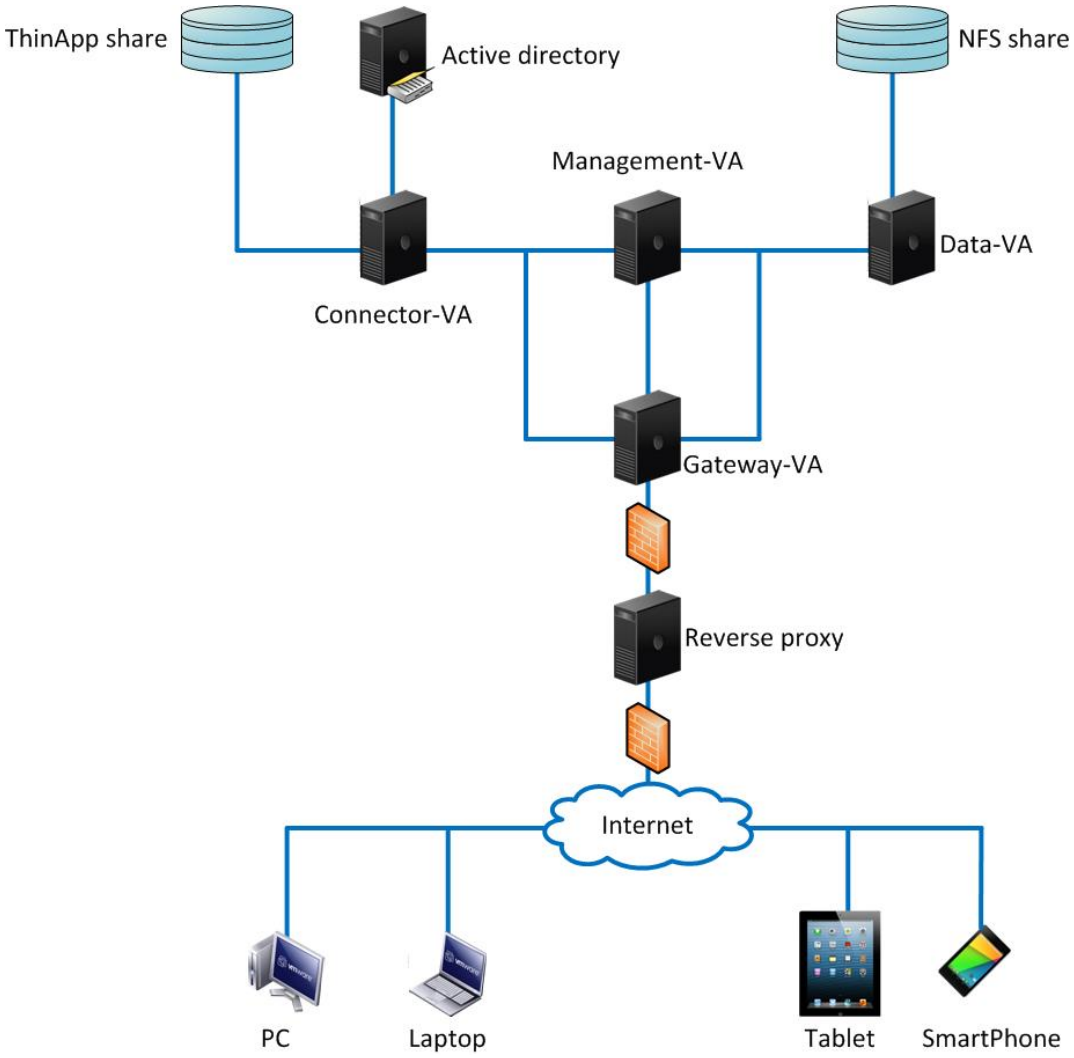

Figure 15. Horizon workspace architecture layout

Using Horizon data with VSPEX architectures

The VSPEX End-User Computing for VMware View environment with added infrastructure supports Horizon Data as depicted in Figure 16. You specify server capacity in generic terms for minimum CPU and memory requirements. The customer is free to select the server and networking hardware that meets or exceeds the stated minimum requirements. The recommended storage delivers a highly available architecture for your Horizon Data deployment.

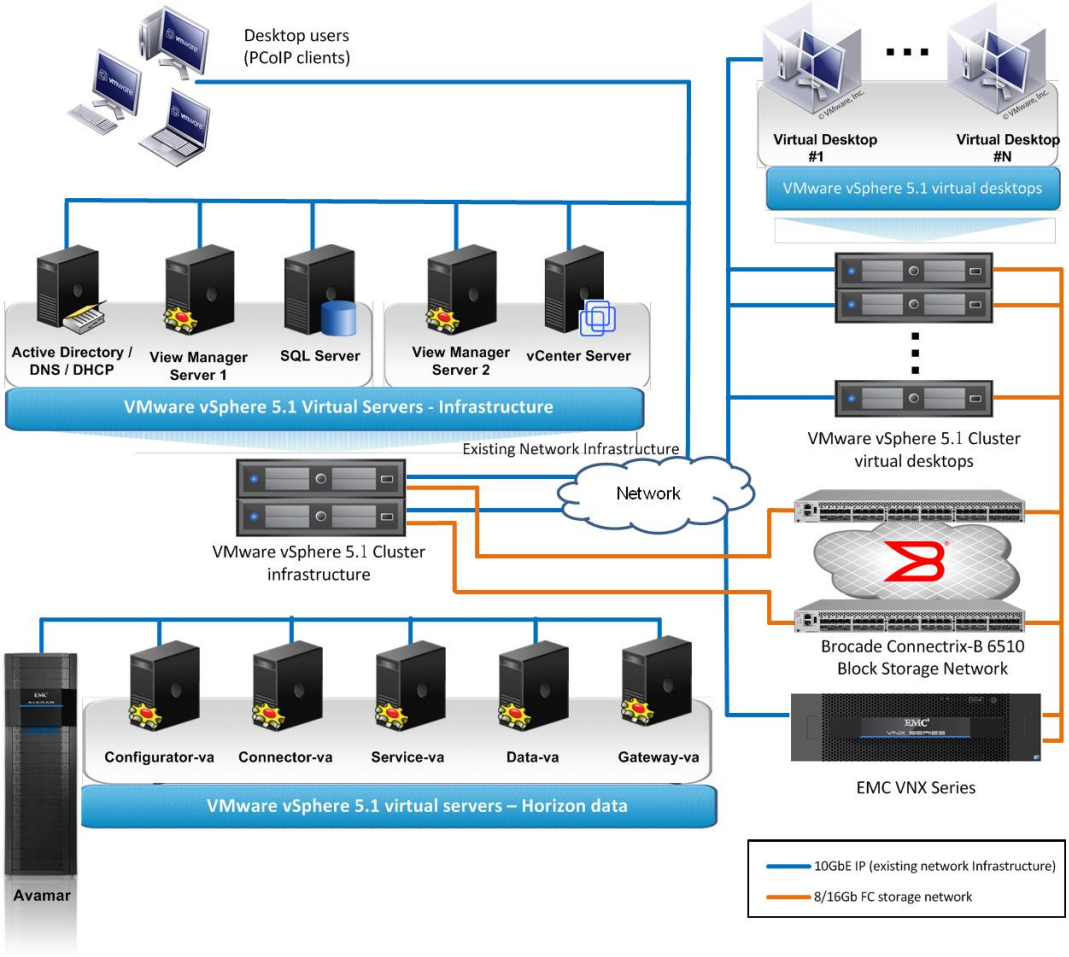

Figure 16. Logical architecture: VSPEX End-User Computing for VMware View with Horizon Data

Server requirements

Table 4 details the minimum supported hardware requirements of each virtual appliance in the VMware Horizon Workspace vApp.

Table 4. Minimum hardware resources to support VMware Horizon data

vApp	vCPU	Memory (GB)	Disks (GB)
Configurator-va	1	1	5
Service-va	2	4	36
Connector-va	2	4	12
Data-va	2	4	350
Gateway-va	1	1	9

Note: For high availability during failure scenarios, it may be necessary to restart virtual machines on different hardware, those physical servers will need to have resources available. Follow the specific recommendations in Compute layer to enable this functionality.

Networking components

The networking components can be implemented using 1 Gb or 10 Gb IP networks, provided that bandwidth and redundancy are sufficient to meet the listed requirements.

Storage component

VMware Horizon Data can use NFS or block storage to provide data service. The EMC VNX storage family is the number one external storage platform in the industry. Its ability to provide file and block access with a broad feature set makes it an ideal choice for VMware Horizon Data.

The EMC VNX series supports a wide range of business class features ideal for Horizon Data including:

- Thin Provisioning
- Fully Automated Storage Tiering for Virtual Pools
- Fast Cache
- Data Compression/file deduplication
- Replication
- Checkpoints
- File-Level Retention
- Quota Management and many more

Table 5. Recommended EMC VNX storage needed for the Horizon Data NFS share

NFS shares for:	Configuration	Notes
500 users:	• Two data movers (active/standby CIFS variant only) • Eight 2 TB, 7,200 rpm 3.5-inch NL-SAS disks	Provided that each user has 10 GB of private storage space
1,000 users	• Two Data Movers (active/standby CIFS variant only) • Sixteen 2 TB, 7,200 rpm 3.5-inch NL-SAS disks	
2,000 users	• Two Data Movers (active/standby CIFS variant only) • Twenty-four 2 TB, 7,200 rpm 3.5-inch NL-SAS disks	

Chapter 4 Solution Architectural Overview

This chapter presents the following topics:

Solution overview 60

Solution architecture 60

Server configuration guidelines 73

Brocade Network configuration guidelines 76

Storage configuration guidelines 81

High availability and failover 90

Validation test profile 94

Antivirus and antimalware platform profile 95

vCenter Operations Manager for View platform profile desktops 96

Backup and recovery configuration guidelines 98

Sizing guidelines 98

Reference workload 98

Applying the reference workload 99

Implementing the reference architectures 100

Quick assessment 103

Solution overview

VSPEX Proven Infrastructure solutions with Brocade networking are built with proven best-of-breed technologies to create a complete virtualization solution that enables you to make an informed decision when choosing and sizing the hypervisor and compute layers. VSPEX eliminates many server virtualization planning and configuration burdens by leveraging extensive interoperability, functional, and performance testing by EMC. VSPEX accelerates your IT Transformation to cloud-based computing by enabling faster deployment, more choice, higher efficiency, and lower risk.

This chapter includes a comprehensive guide to the major aspects of this solution. The customer is free to select the server hardware of their choice that meets or exceeds the stated minimums; server capacity is specified in generic terms for required minimums of CPU, memory, and storage network with Brocade Fabric network switch interfaces. The specified storage architecture, along with a system meeting the server and Brocade storage network requirements outlined, has been validated by EMC to provide high levels of performance while delivering a highly available architecture for your End-User Computing deployment.

Each VSPEX Proven Infrastructure balances the storage, network, and compute resources needed for a set number of virtual desktops, which have been validated by EMC. In practice, each virtual machine has its own set of requirements that rarely fit a pre-defined idea of what a virtual machine should be. In any discussion about virtual infrastructures, it is important to first define a reference workload. Not all servers perform the same tasks, and it is impractical to build a reference that takes into account every possible combination of workload characteristics.

Solution architecture

The VSPEX End-User-Computing solution for up to 2,000 virtual desktops is validated at three different points of scale. These defined configurations form the basis of creating a custom solution. These points of scale are defined in terms of the reference workload.

Note: VSPEX uses the concept of a Reference Workload to describe and define a virtual machine. Therefore, one physical or virtual desktop in an existing environment may not be equal to one virtual desktop in a VSPEX solution. Evaluate your workload in terms of the reference to arrive at an appropriate point of scale. The detailed process is described in Applying the reference workload.

Logical architecture

The architecture diagrams in this section show the layout of the major components comprising the solutions. This solution includes two variants: block storage and file storage.

Two networks are in use: one, Brocade storage network for carrying virtual desktop and virtual server operating system (OS) data and one 10 Gb Ethernet for carrying all other traffic. The Brocade storage network can use 8 or 16 Gb FC, 10 Gb Ethernet with FCoE, or 10 Gb Ethernet with iSCSI protocol.

Figure 17 shows the logical architecture of implementation in block storage.

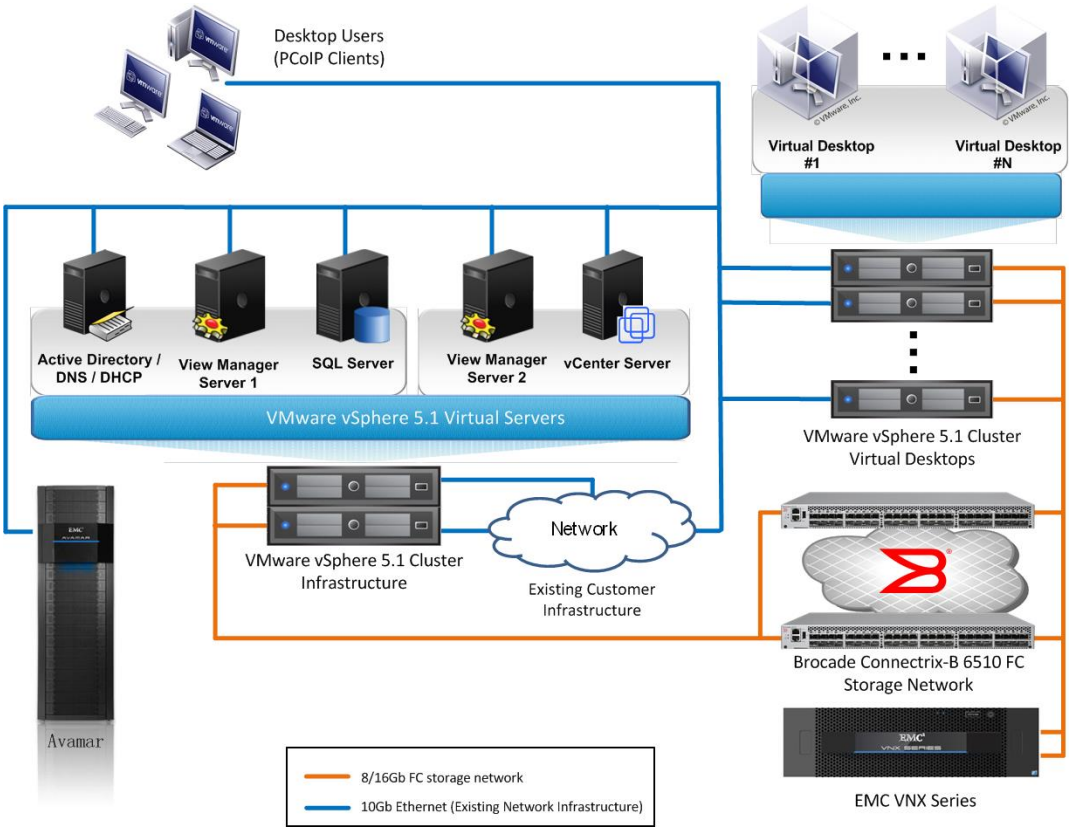

Figure 17. Logical architecture for block storage

Note: This solution also supports 1 Gb Ethernet if the bandwidth requirements are met.

Figure 18 shows the file storage logical architecture. The 10 GbE IP network carries all traffic.

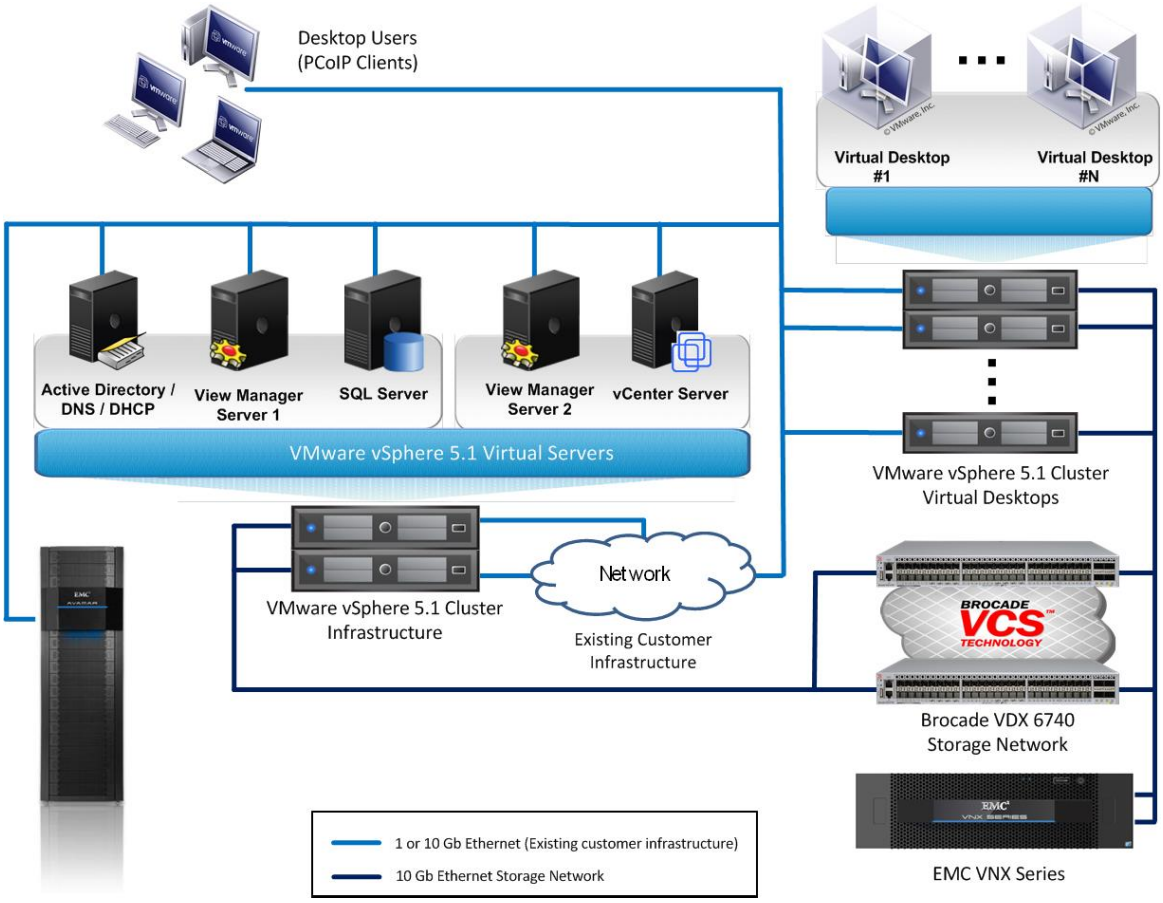

Figure 18. Logical architecture for NFS storage

Note: This solution also supports 1 Gb Ethernet if the bandwidth requirements are met.

Key components

VMware Horizon View Manager Server 5.3—Provides virtual desktop delivery, authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops. In this solution architecture, VMware Horizon View Manager 5.3 is installed on Windows Server 2008 R2 and hosted as a virtual machine on a VMware vSphere server. Two VMware Horizon View Manager Servers are used in this solution.

Virtual desktops—Persistent virtual desktops running Windows 7 are provisioned as VMware Horizon View Linked Clones.

VMware vSphere—Provides a common virtualization layer to host a server environment that contains the virtual machines. The specifics of the validated environment are listed in Table 6. VSphere provides a highly available infrastructure through such features as:

- **vMotion**—Provides live migration of virtual machines within a virtual infrastructure cluster with no virtual machine downtime or service disruption.
- **Storage vMotion**—Provides live migration of virtual machine disk files within and across storage arrays with no virtual machine downtime or service disruption.
- **vSphere High Availability (HA)**—Detects and provides rapid recovery for a failed virtual machine in a cluster.
- **Distributed Resource Scheduler (DRS)**—Provides load balancing of computing capacity in a cluster.
- **Storage Distributed Resource Scheduler (SDRS)**—Provides load balancing across multiple datastores, based on space use and I/O latency.

VMware vCenter Server—Provides a scalable and extensible platform that forms the foundation for virtualization management for the VMware vSphere cluster. All vSphere hosts and their virtual machines are managed through vCenter.

VMware vShield Endpoint—Offloads virtual desktop antivirus and antimalware scanning operations to a dedicated secure virtual appliance delivered by VMware partners. Offloading scanning operations improves desktop consolidation ratios and performance by eliminating antivirus storms, streamlining antivirus and antimalware deployment, and monitoring and satisfying compliance and audit requirements through detailed antivirus and antimalware activity logging.

VMware vCenter Operations Manager for View (vCOps)— monitors the virtual desktops and all of the supporting elements of the VMware Horizon View virtual infrastructure.

VSI for VMware vSphere—A plug-in to the vSphere Client that provides storage management for EMC arrays directly from the Client. VSI is highly customizable and helps provide a unified management interface.

VMware vCenter SQL Server—Requires a database service to store configuration and monitoring details. A Microsoft SQL Server 2008 R2 running on a Windows 2008 R2 Server is used for this purpose.

DHCP server—Centrally manages the IP address scheme for the virtual desktops. This service is hosted on the same virtual machine as the domain controller and DNS server. The Microsoft DHCP Service running on a Windows 2012 server is used for this purpose.

DNS Server—Required for the various solution components to perform name resolution. The Microsoft DNS Service running on a Windows Server 2012 server is used for this purpose.

Active Directory Server—Services that are required for the various solution components to function properly. The Microsoft AD Directory Service running on a Windows Server 2012 server is used for this purpose.

Shared Infrastructure —DNS and authentication/authorization services like Microsoft Active Directory can be provided using the existing infrastructure or set up as part of the new virtual infrastructure.

IP network—A shared IP network carries user and management traffic. A standard Ethernet network should use redundant cabling and switching.

Brocade Storage Network — VSPEX with Brocade networking offers different options for block-based and file-based storage networks. All storage traffic is carried over redundant cabling and Brocade Fabric switches.

Storage Network for Block:

This solution provides three options for block based storage networks.

- **Fibre Channel (FC)** is a set of standards that define protocols for performing high speed serial data transfer. FC provides a standard data transport frame among servers and shared storage devices.
- **Brocade 6510 Fibre Channel Switch** — Provides fast, easy and resilience scaling from 24 to 48 Ports on Demand (PoD) capabilities and supports 2,4, 8, or 16 Gbps speeds for FC attached VNX5400, VNX5600, and VNX5800 arrays .
- **Fibre Channel over Ethernet (FCoE)** is a new storage networking protocol that supports FC natively over Ethernet, by encapsulating FC frames into Ethernet frames. This allows the encapsulated FC frames to run alongside traditional Internet Protocol (IP) traffic.
- **Brocade VDX 6740 Ethernet Fabric Switch** — Provides efficient, easy to configure, resiliency that scales from 24 to 64 Port on Demand (PoD) at 10GbE for FCoE attached VNX5400, VNX 5600 and VNX 5800 arrays.
- **10 Gb Ethernet (iSCSI)** enables the transport of SCSI blocks over a TCP/IP network. iSCSI works by encapsulating SCSI commands into TCP packets and sending the packets over the IP network.
- **Brocade VDX 6740 Ethernet Fabric Switch** — Provides efficient, easy to configure, resiliency that scales from 24 to 64 Port on Demand (PoD) at 1 GbE or 10GbE for iSCSI attached VNX5400, VNX 5600 and VNX 5800 arrays.

Storage Network for File:

With file-based storage, a10Gb Ethernet private network carries the storage traffic. A Brocade **10 Gb Ethernet Fabric network** enables the transport of File for NFS and CIFS storage network.

- **Brocade VDX 6740 Ethernet Fabric Switch** — Provides efficient, easy to configure, resiliency that scales from 24 to 64 Port on Demand (PoD) at 1 GbE or 10GbE for file attached VNX5400, VNX 5600 and VNX 5800 arrays.

EMC VNX5400 array—Provides storage by presenting NFS/FC datastores to vSphere hosts for up to 1,000 virtual desktops.

EMC VNX5600 array — Provides storage by presenting NFS/FC datastores to vSphere hosts for up to 2,000 virtual desktops.

VNX family storage arrays

VNX family storage arrays include the following components:

- **Storage processors (SPs)** support block data with UltraFlex I/O technology that supports Fibre Channel, iSCSI, and FCoE protocols. The SPs provide access for all external hosts and for the file side of the VNX array.
- The **disk-processor enclosure (DPE)** is 3U in size and houses each storage processor as well as the first tray of disks. This form factor is used in the VNX5400 and VNX5600.
- **X-Blades** (or Data Movers) access data from the back end and provide host access using the same UltraFlex I/O technology that supports the NFS, CIFS, MPFS, and pNFS protocols. The X-Blades in each array are scalable and provide redundancy to ensure that no single point of failure exists.
- The **Data Mover enclosure (DME)** is 2U in size and houses the Data Movers (X-Blades). The DME is similar, in form, to the SPE and is used on all VNX models that support file systems.
- **Standby power supplies** are 1U in size and provide enough power to each storage processor to ensure that any data in flight is de-staged to the vault area in the event of a power failure. This ensures that no writes are lost. Upon restart of the array, the pending writes are reconciled and persisted.
- **Control Stations** are 1U in size and provide management functions to the file-side components referred to as X-Blades. The Control Station is responsible for X-Blade failover. The Control Station may optionally be configured with a matching secondary Control Station to ensure redundancy on the VNX array.
- **Disk-array enclosures (DAE)** house the drives used in the array.

Hardware resources

Table 6 lists the hardware used in this solution.

Table 6. Solution hardware

Hardware	Configuration	Notes
Servers for virtual desktops	CPU: • 1 vCPU per desktop (8 desktops per core) • 63 cores across all servers for 500 virtual desktops • 125 cores across all servers for 1,000 virtual desktops • 250 cores across all servers for 2,000 virtual desktops Memory: • 2 GB RAM per virtual desktop • 1 TB RAM across all servers for 500 virtual desktops • 2 TB RAM across all servers for 1,000 virtual desktops • 4 TB RAM across all servers for 2,000 virtual machines • 2 GB RAM reservation per vSphere host Network: • Three 10 GbE NICs per blade chassis or two 10 GbE NICs per standalone server **Note:** To implement VMware vSphere High Availability (HA) functionality and to meet the listed minimums, the infrastructure should have one additional server.	Add CPU and RAM as needed for the VMware vShield Endpoint and Avamar components. Refer to the vendor documentation for specific details concerning vShield Endpoint and Avamar resource requirements.
Brocade NFS and CIFS storage network infrastructure	Two Brocade VDX 6740 Ethernet Fabric switches – (minimum switching capacity): • Two 10 GbE ports per storage processor for data • Two 10 GbE ports per server or • Six 1GbE NICs per server (optional) **Note:** One 1 GbE port per storage processor for management traffic.	Redundant Brocade VDX Ethernet Fabric configuration

Hardware	Configuration	Notes
Brocade FC-Block storage network infrastructure	Two Brocade Connectrix-B Fibre Channel Switches (minimum switching capability): • Four 4/8 Gb FC ports, or four 10 Gb CEE ports, or four 10 Gb Ethernet ports for VNX backend • Two 8/16 Gb FC ports per storage processor (block only) **Note:** To implement FCoE or iSCSI Block storage network, use Brocade VDX 6740 Ethernet Fabric switches (10GbE only).	Redundant LAN/SAN configuration
Storage	Common • 2 x 10 GbE interfaces per data mover • 1 x 1 GbE interface per control station for management • 2 x 8 Gb FC ports per storage processor (FC only) For 500 virtual desktops • 2 x Data Movers (active/standby) • 15 x 300 GB, 15 k rpm 3.5-inch SAS disks • 3 x 100 GB, 3.5-inch flash drives For 1,000 virtual desktops • 2 x Data Movers (active/standby) • 20 x 300 GB, 15 k rpm 3.5-inch SAS disks • 3 x 100 GB, 3.5-inch flash drives For 2,000 virtual desktops • 3 x Data Movers (Two active, one standby) • 36 x 300 GB, 15 k rpm 3.5-inch SAS disks • 5 x 100 GB, 3.5-inch flash drives	VNX shared storage
	For 500 virtual desktops, • 9 x 2 TB, 7,200 rpm 3.5-inch NL-SAS disks For 1,000 virtual desktops • 17 x 2 TB, 7,200 rpm 3.5-inch NL-SAS disks For 2,000 virtual desktops • 34 x 2 TB, 7,200 rpm 3.5-inch NL-SAS disks	Optional for user data

Hardware	Configuration	Notes
	For 500 virtual desktops • 5 x 300 GB, 15k rpm 3.5-inch SAS disks For 1,000 virtual desktops • 5 x 300 GB, 15 k rpm 3.5-inch SAS disks For 2,000 virtual desktops • 10 x 300 GB, 15k rpm 3.5-inch SAS disks	Optional for infrastructure storage
	For 500 virtual desktops • 5 x 300 GB, 15k rpm 3.5-inch SAS disks For 1,000 virtual desktops • 5 x 300 GB, 15 k rpm 3.5-inch SAS disks For 2,000 virtual desktops • 10 x 300 GB, 15k rpm 3.5-inch SAS disks	Optional for vCenter Operations Manager for View
Servers for customer infrastructure	Minimum number required for 500 virtual desktops: • 2 physical servers • 20 GB RAM per server • Eight processor cores per server • Four 1 GbE ports per server • Additional CPU and RAM as needed for the VMware vShield Endpoint components. Minimum number required for 1,000 virtual desktops: • 2 physical servers • 48 GB RAM per server • Eight processor cores per server • Four 1 GbE ports per server • Additional CPU and RAM as needed for the VMware vShield Endpoint components. Minimum number required for 2,000 virtual desktops: • 2 physical servers • 48 GB RAM per server • Eight processor cores per server • Four 1 GbE ports per server • Additional CPU and RAM as needed for the VMware vShield Endpoint components.	These servers and the roles they fulfill may already exist in the customer environment Refer to vendor documentation for specific details concerning vShield Endpoint resource requirements
EMC backup	See *EMC backup and recovery options for VSPEX End User Computing with VMware Horizon View* available on EMC Online Support.	

Software resources

Table 7 lists the software used in this solution.

Table 7. Solution software

Software	Configuration
VNX5400/5600 (shared storage, file systems)	
VNX OE for file	Release 8.0
VNX OE for block	Release 33 (05.33)
EMC VSI for VMware vSphere: Unified Storage Management	VSI 5.6
EMC VSI for VMware vSphere: Storage Viewer	VSI 5.6
VMware Horizon View Desktop Virtualization	
VMware Horizon View Manager Server	Version 5.3 Premier
Operating system for VMware Horizon View Manager	Windows Server 2008 R2 Standard Edition
Microsoft SQL Server	Version 2008 R2 Standard Edition
EMC Avamar next-generation backup	
Avamar	7.0
Avamar Agent	7.0
VMware vSphere	
vSphere Server	5.5*
vCenter Server	5.5.
vShield Manager (includes vShield Endpoint Service)	5.5
Operating system for vCenter Server	Windows Server 2008 R2 Standard Edition
vStorage API for Array Integration Plug-in (VAAI) (NFS variant only)	1.0-11
PowerPath Virtual Edition (FC variant only)	PowerPath/VE 5.9
VMware vCenter Operations Manager for View	
VMware vCenter Operations Manager	5.7.0
vCenter Operations Manager for View plug-in	1.5

Software	Configuration
Virtual Desktops **Note** Aside from the base operating system, this software was used for solution validation and is not required	
Base operating system	Microsoft Windows 7 Enterprise (32-bit) SP1
Microsoft Office	Office Enterprise 2007 Version 12
Internet Explorer	8.0.7601.17514
Adobe Reader	X (10.1.3)
VMware vShield Endpoint (component of VMware Tools)	9.0.5 build-1065307
Adobe Flash Player	11
Bullzip PDF Printer	7.2.0.1304
FreeMind	0.8.1
Login VSI (VDI workload generator)	3.7 Professional Edition
* Patch ESXi510-201210001 needed for support View 5.3	

Sizing for validated configuration

When selecting servers for this solution, the processor core shall meet or exceed the performance of the Intel Nehalem family at 2.66 GHz. As servers with greater processor speeds, better performance, and higher core density become available, servers can be consolidated as long as the required total core and memory count is met and a sufficient number of servers are incorporated to support the necessary level of high availability.

As with servers, network interface card (NIC) speed and quantity can also be consolidated as long as the overall bandwidth requirements for this solution and sufficient redundancy necessary to support high availability are maintained.

Table 8 represents a sample server configuration required to support specific desktop solutions.

Table 8. Sample server configuration

Storage	Configuration		
Number of desktops	500	1000	2000
Number of servers	8	16	32
Processor type	Intel Nehalem 2.66 GHz		
Processor per server	Two 4-core processors		
RAM (GB) per server	128		
NIC	One 10 GbE per every four-blade server plus one 10 GbE for each blade chassis		

As shown in Table 6, a minimum of one core is required to support eight virtual desktops with a minimum of 2 GB of RAM for each. You must also consider the correct balance of memory and cores for the expected number of virtual desktops to be supported by a server. Additional CPU resources and RAM will be required to support the VMware vShield Endpoint components. Read the vendor documentation for specific details.

The Brocade VDX Ethernet Fabric switches deployed in this solution architecture for the storage network exceed the minimum backplane capacity of 96 (for 500 virtual desktops), 192 (for 1,000 virtual desktops) or 320 (for 2,000 virtual desktops) Gb/s non-blocking required and supports the following features for the storage network in the VSPEX architectures:

- IEEE 802.1x Ethernet flow control
- 802.1q VLAN tagging
- Ethernet link aggregation using IEEE 802.1ax (802.3ad) Link Aggregation Control Protocol
- Simple Network Management Protocol (SNMP) management capability
- Jumbo frames

The Brocade network switches with the VSPEX solutions supports scalable bandwidth and high availability. Brocade network deployed with the components of a VSPEX End User Computing solution *provides a storage network configuration* that provides the following:

- Two switch deployment to support redundancy
- Redundant power supplies
- Scalable port density for a minimum of forty 1 GbE or eight 10 GbE ports (for 500 virtual desktops), two 1 GbE and sixteen 10 GbE ports (for 1,000 virtual desktops) or two 1 GbE and thirty-two 10 GbE ports (for 2,000 virtual desktops), distributed for high availability.
- The appropriate uplink ports for customer connectivity

Use of 10 GbE ports should align with those on the server and storage while keeping in mind the overall network requirements for this solution and a level of redundancy to support high availability. Additional server NICs and storage connections should also be considered based on customer or specific implementation requirements.

Note: For servers and VNX series array that are connect with FC based block storage network, use the Brocade Connectrix-B Fibre Channel switches. The Brocade Connectrix-B FC switches provide the 8 or 16 Gb ports and exceed the minimum backplane capacity of 96 (for 500 virtual desktops), 192 (for 1,000 virtual desktops) or 320 (for 2,000 virtual desktops) Gb/s non-blocking required. (See the configuration procedure described in Chapter 5.)

The management infrastructure (Active Directory, DNS, DHCP, and SQL Server) can be supported on two servers similar to those previously defined, but will require a minimum of only 48 GB RAM instead of 128 GB.

Disk layout details are provided in Storage configuration guidelines on page 81.

Server configuration guidelines

Overview

When designing and ordering the compute/server layer of the VSPEX solution described below, several factors might alter the final purchase. From a virtualization perspective, if a system's workload is well understood, features like Memory Ballooning and Transparent Page Sharing can reduce the aggregate memory requirement.

If the virtual machine pool does not have a high level of peak or concurrent usage, the number of vCPUs can be reduced. Conversely, if the applications being deployed are highly computational in nature, the number of CPUs and memory purchased might need to be increased. Table 9 identifies the server hardware and the configurations used in this solution.

Table 9. Server hardware

Servers for virtual desktops	Configuration
CPU:	• 1 vCPU per desktop (8 desktops per core) • 63 cores across all servers for 500 virtual desktops • 125 cores across all servers for 1000 virtual desktops • 250 cores across all servers for 2000 virtual desktops
Memory:	• 2 GB RAM per virtual machine • 1 TB RAM across all servers for 500 virtual desktops • 2 TB RAM across all servers for 1,000 virtual desktops • 4 TB RAM across all servers for 2,000 virtual machines • 2 GB RAM reservation per vSphere host
Network	• 6 x 1 GbE NICs per server for 500 virtual desktops • 3 x 10 GbE NICs per blade chassis or 6 x 1 GbE NICs per standalone server for 1,000 virtual desktops • 3 x 10 GbE NICs per blade chassis or 6 x 1 GbE NICs per standalone server for 2,000 virtual desktops

Notes:

- Add CPU and RAM as needed for the VMware vShield Endpoint and Avamar components. Refer to the vendor documentation for details concerning vShield Endpoint and Avamar resource requirements.
- To implement VMware vSphere High Availability (HA) and to meet the listed minimum requirements, the infrastructure should have one additional server.

vSphere memory virtualization for VSPEX

VMware vSphere has a number of advanced features that help to maximize performance and overall use of resources. This section describes some of the important features that help to manage memory and considerations for using them in the environment.

In general, you can consider virtual machines on a single hypervisor consuming memory as a pool of resources:

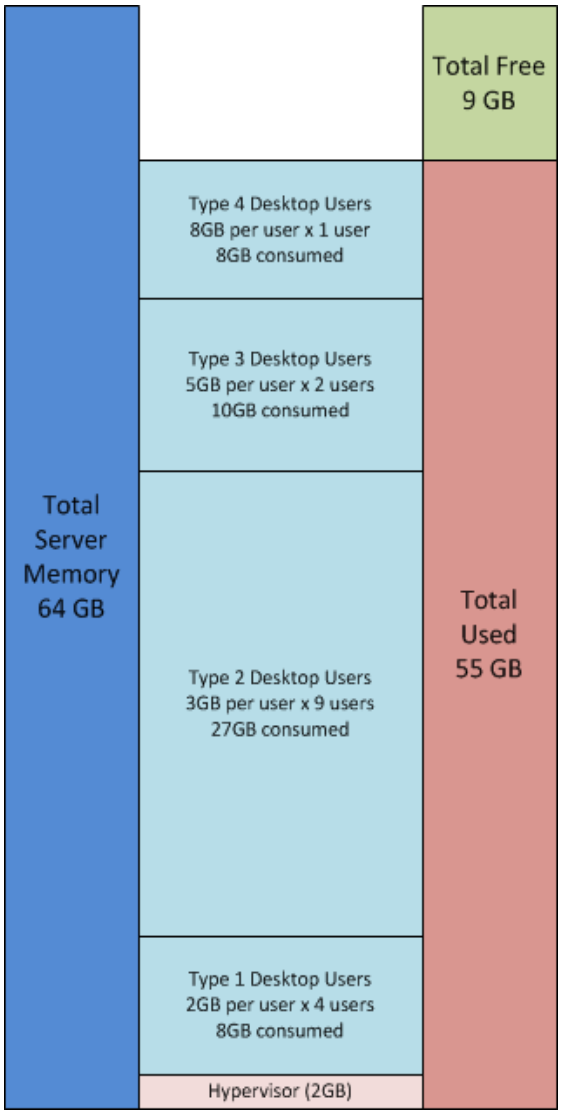

Figure 19. Hypervisor memory consumption

You should understand the technologies presented in this section.

Memory over-commitment

Memory over-commitment occurs when more memory is allocated to virtual machines than is physically present in a VMware vSphere host. Using sophisticated techniques, such as ballooning and transparent page sharing, vSphere is able to handle memory over-commitment without any performance degradation. However, if more memory than is present on

the server is being actively used, vSphere might resort to swapping portions of a virtual machine's memory.

Non-Uniform Memory Access (NUMA)

vSphere uses a NUMA load-balancer to assign a home node to a virtual machine. Memory access is local and provides the best performance possible because memory for the virtual machine is allocated from the home node. Applications that do not directly support NUMA benefit from this feature.

Transparent page sharing

Virtual machines running similar operating systems and applications typically have identical sets of memory content. Page sharing allows the hypervisor to reclaim the redundant copies and keep only one copy, which frees up the total host memory consumption. If most of your application virtual machines run the same operating system and application binaries, the total memory usage can be reduced to increase consolidation ratios.

Memory ballooning

By using a balloon driver loaded in the guest operating system, the hypervisor can reclaim host physical memory if memory resources are under contention with little to no impact on the application's performance.

Memory configuration guidelines

This section provides guidelines for allocating memory to virtual machines, which take into account vSphere memory overhead and the virtual machine memory settings.

vSphere memory overhead

There is some associated overhead for the virtualization of memory resources. The memory space overhead has two components.

- The system overhead for the VMkernel
- Additional overhead for each virtual machine

The amount of additional overhead memory for the VMkernel is fixed, whereas the amount of additional memory for each virtual machine depends on the number of virtual CPUs and the configured memory for the guest operating system.

Allocating memory to virtual machines

The proper sizing of memory for a virtual machine in VSPEX architectures is based on many factors. Because of the number of application services and use cases available, determining a suitable configuration for an environment requires creating a baseline configuration, testing, and making adjustments, as discussed in this paper. In this solution, each virtual machine is assigned two GB of memory in fixed mode, as listed in Table 6.

Brocade Network configuration guidelines

Overview

This section provides guidelines for setting up a redundant, highly available storage network configuration. The guidelines outlines compute access to existing infrastructure, management network and Brocade storage network for compute to EMC unified storage.

Administrators use the Management Network as a dedicated way to access the management connections on the storage array, network switches, and hosts. The Storage Network is communication between the compute layer and the storage layer. The Brocade Storage Network provides the communication between the compute layer and the storage layer. Storage network guidelines are outlined for configuring block and file with the VNX unified storage: File based storage network connectivity with Jumbo Frames, Link Aggregation Control Protocol (LACP), and VLAN features; Block based storage network with 8 Gbps Fibre Channel connectivity with zoning configuration guidelines. For detailed Brocade storage network resource requirements, refer to Table 6.

Table 10. Hardware resources for network infrastructure

Minimum switching capacity for	Configuration
Block	**Brocade Fibre Channel Switch** Two **Brocade 6510 switches** – 24 to 48 Ports on Demand(PoD) • 2 x FC ports per VMware vSphere server, for storage network • 2 x FC ports per SP, for desktop data (Use Brocade VDX 6740 Ethernet Fabric switch for FCoE and iSCSI 10GbE ports on VNX SP) **Note:** Use existing network infrastructure for data and management traffic. • 2 x 10 GbE ports per VMware vSphere server • 2 x 10 GbE ports per Data Mover for user data • 1 x 1 GbE port per Control Station for management
File	**Brocade VDX Ethernet Fabric Switch** Two **Brocade VDX 6740 switches** – 24 to 64 PoD • 4 x 10 GbE ports per VMware vSphere server • 2 x 10 GbE ports per Data Mover for data **Note:** Use existing network infrastructure for data and management traffic. • 1 x 1 GbE port per Control Station for management

Notes:

- The solution may use 1 Gb network infrastructure as long as the underlying requirements around bandwidth and redundancy are fulfilled.
- This table assumes that the VSPEX implementation is using rack mounted servers; for implementations based on blade servers, ensure that similar bandwidth and high availability capabilities are available.

Enable jumbo frames (for iSCSI and NFS)

Brocade VDX Series switches support the transport of jumbo frames. This solution for EMC VSPEX private cloud recommends an MTU set at 9216 (jumbo frames) for efficient storage and migration traffic. Jumbo frames are enabled by default on the Brocade ISL trunks. However, to accommodate end-to-end jumbo frame support on the network for the edge hosts, this feature can be enabled under the vLAG interface connected to the ESXi hosts, and the VNX NFS server. The default Maximum Transmission Unit (MTU) on these interfaces is 2500. This MTU is set to 9216 to optimize the network for jumbo frame support.

Link Aggregation

A link aggregation resembles an Ethernet channel, but uses the Link Aggregation Control Protocol (LACP) IEEE 802.3ad standard. The IEEE 802.3ad standard supports link aggregations with two or more ports. All ports in the aggregation must have the same speed and be full duplex. In this solution, Link Aggregation Control Protocol (LACP) is configured on VNX, combining multiple Ethernet ports into a single virtual device. If a link is lost in the Ethernet port, the link fails over to another port. All network traffic is distributed across the active links.

Brocade Virtual Link Aggregation Group (vLAG)

Brocade Virtual Link Aggregation Groups (vLAGs) are used for the ESXi hosts, the VNX array, and the VMware NFS server. In the case of the VNX, a dynamic Link Aggregation Control Protocol (LACP) vLAG is used. In the case of ESXi hosts static LACP vLAGs are used. While Brocade ISLs are used as interconnects between Brocade VDX switches within a Brocade VCS fabric, industry standard LACP LAGs are supported for connecting to other network devices outside the Brocade VCS fabric. Typically, LACP LAGs can only be created using ports from a single physical switch to a second physical switch. In a Brocade VCS fabric, a vLAG can be created using ports from two Brocade VDX switches to a device to which both VDX switches are connected. This provides an additional degree of device-level redundancy, while providing active-active link-level load balancing.

Brocade Inter-Switch Link (ISL) Trunks

In VSPEX Stack Brocade Inter-Switch Link (ISL) Trunking is used within the Brocade VCS fabric to provide additional redundancy and load balancing between the NFS clients and NFS server. Typically, multiple links between two switches are bundled together in a Link Aggregation Group (LAG) to provide redundancy and load balancing. Setting up a LAG requires lines of configuration on the switches and selecting a hash-based algorithm for load balancing based on source-destination IP or MAC addresses.

All flows with the same hash traverse the same link, regardless of the total number of links in a LAG. This might result in some links within a LAG, such as those carrying flows to a storage target, being over utilized and packets being dropped, while other links in the LAG remain underutilized. Instead of LAG-based switch interconnects, Brocade VCS Ethernet fabrics automatically form ISL trunks when multiple connections are added between two Brocade VDX® switches. Simply adding another cable increases bandwidth, providing linear scalability of switch-to-switch traffic, and this does not require any configuration on the switch. In addition, ISL trunks use a frame-by-frame load balancing technique, which evenly balances traffic across all members of the ISL trunk group.

Equal-Cost Multipath (ECMP)

A standard link-state routing protocol that runs at Layer 2 determines if there are Equal-Cost Multipaths (ECMPs) between RBridges in an Ethernet fabric and load balances the traffic to make use of all available ECMPs. If a neighbor switch is reachable via several interfaces with different bandwidths, all of them are treated as "equal-cost" paths. While it is possible to set the link cost based on the link speed, such an algorithm complicates the operation of the fabric. Simplicity is a key value of Brocade VCS Fabric technology, so an implementation is chosen in the test case that does not consider the bandwidth of the interface when selecting equal-cost paths. This is a key feature needed to expand network capacity, to keep ahead of customer bandwidth requirements.

Pause Flow Control

Pause Flow Control is enabled on vLAG-facing interfaces connected to the ESXi hosts, and the NFS server. Brocade VDX Series switches support the Pause Flow Control feature. IEEE 802.3x Ethernet pause and Ethernet Priority-Based Flow Control (PFC) are used to prevent dropped frames by slowing traffic at the source end of a link. When a port on a switch or host is not ready to receive more traffic from the source, perhaps due to congestion, it sends pause frames to the source to pause the traffic flow. When the congestion is cleared, the port stops requesting the source to pause traffic flow, and traffic resumes without any frame drop. When Ethernet pause is enabled, pause frames are sent to the traffic source. Similarly, when PFC is enabled, there is no frame drop; pause frames are sent to the source switch.

VLAN

VLANs isolate network traffic to allow the traffic between hosts and storage, hosts and clients, and management traffic to move over isolated networks. In some cases, physical isolation may be required for regulatory or policy compliance reasons; in many cases, logical isolation using VLANs is sufficient. This solution calls for a minimum of three VLANs:

- Client access
- Storage
- Management

Figure 20 illustrates the VLANs.

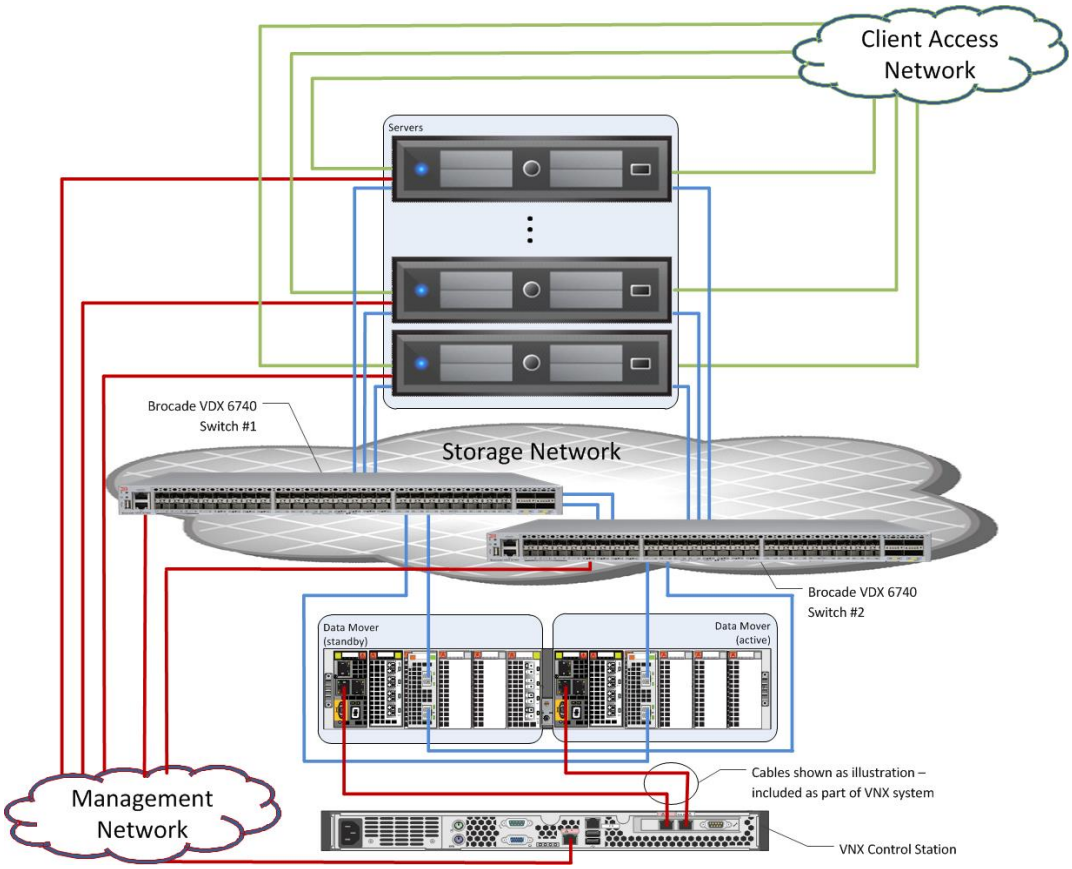

Figure 20. Required Brocade VDX network

Note: The diagram demonstrates the network connectivity requirements for a VNX array using 10 GbE network connections. A similar topology should be created when using 1 GbE network connections.

The client access network is for users of the system, or clients, to communicate with the infrastructure. The Storage Network is used for communication between the compute layer and the storage layer. The Management network is used for administrators to have a dedicated way to access the management connections on the storage array, network switches, and hosts.

Notes:
Some best practices call for additional network isolation for cluster traffic, virtualization layer communication, and other features. These additional networks can be implemented but they are not required.
If the Fibre Channel storage network option is picked for the deployment, similar best practices and design principles apply.

Zoning (FC Block Storage Network only)

Zoning is mechanism used to specify the devices in the fabric that should be allowed to communicate with each other for storage network traffic between host & storage (Block based only). Zoning is based on either port World Wide Name (pWWN) or Domain, Port (D, P). (See the Secure SAN Zoning Best Practices white paper in Appendix C for details.) When using pWWN, the SAN administrators cannot pre-provision zone assignments until the servers are connected and the WWN name of the HBAs is known. The Brocade fabric-based implementation supports a scalable solution for environments with blade and rack servers. This solution calls for a minimum of 2 zones:

- Block Storage Network

Figure 21 depicts the VLANs for Client Access network & Management network and the zones for Bock Storage network connectivity requirements for a block-based VNX array.

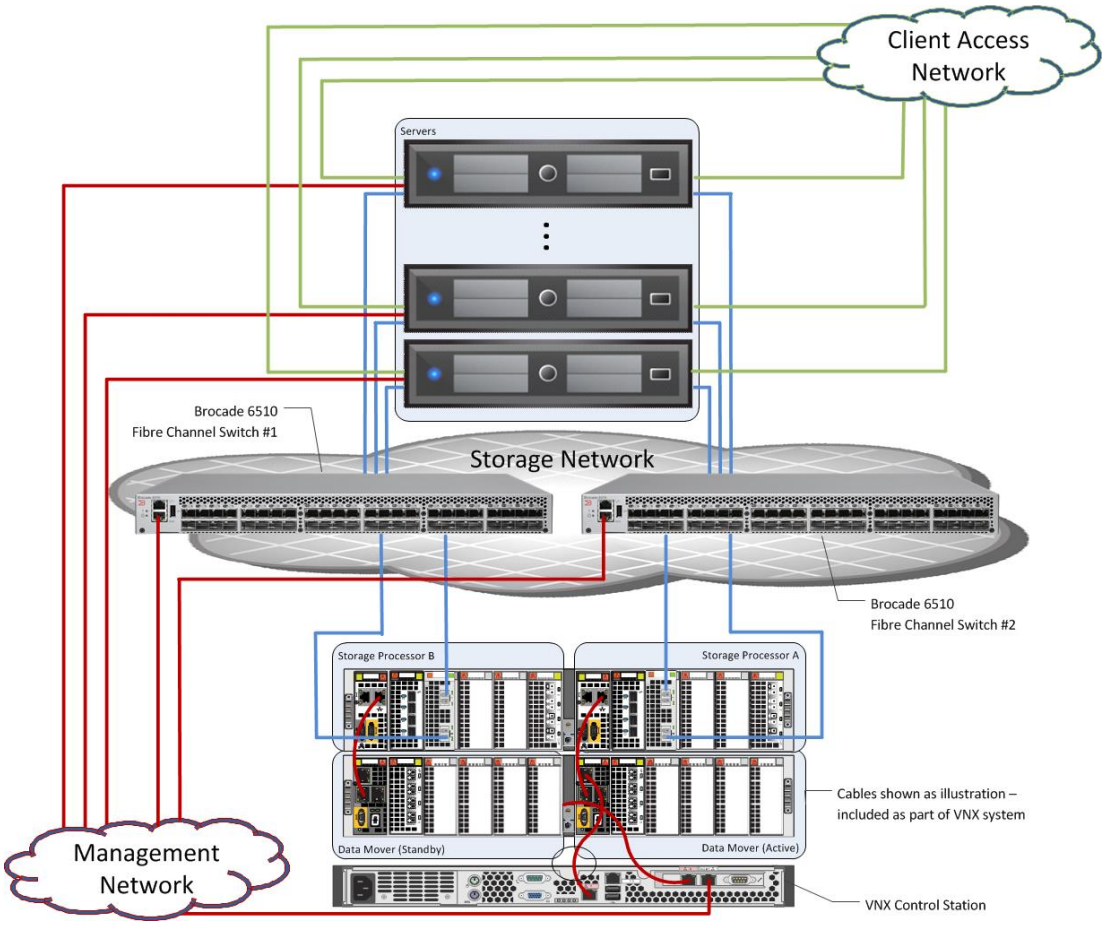

Figure 21. Required networks with block storage variant

Storage configuration guidelines

Overview

vSphere allows more than one method of using storage when hosting virtual machines. The solutions described in Table 11 were tested using NFS or FC, and the storage layout described adheres to all current best practices. An educated customer or architect can make modifications based on their understanding of the system's usage and load if required.

Table 11. Storage hardware

Hardware	Configuration	Notes
Storage	Common • 2 x 10 GbE interfaces per data mover • 1 x 1 GbE interface per control station for management • 2 x 8 Gb FC ports per storage processor (FC only) For 500 virtual desktops • 2 x Data Movers (active/standby) • 15 x 300 GB 15 k rpm 3.5-inch SAS disks • 3 x 100 GB 3.5-inch flash drives For 1,000 virtual desktops • 2 x Data Movers (active/standby) • 20 x 300 GB 15 k rpm 3.5-inch SAS disks • 3 x 100 G, 3.5-inch flash drives For 2,000 virtual desktops • 3 x Data Movers (two active one standby) • 36 x 300 GB 15 k rpm 3.5-inch SAS disks • 5 x 100 GB 3.5-inch flash drives	VNX shared storage
	For 500 virtual desktops • 9 x 2 TB 7200 rpm 3.5-inch NL-SAS disks For 1,000 virtual desktops • 17 x 2 TB 7200 rpm 3.5-inch NL-SAS disks For 2,000 virtual desktops • 34 x 2 TB 7200 rpm 3.5-inch NL-SAS disks	Optional for user data

Hardware	Configuration	Notes
	For 500 virtual desktops • 5 x 300 GB 15 k rpm 3.5-inch SAS disks For 1,000 virtual desktops • 5 x 300 GB 15 k rpm 3.5-inch SAS disks For 2,000 virtual desktops • 10 x 300 GB 15 k rpm 3.5-inch SAS disks	Optional for infrastructure storage
	For 500 virtual desktops • 5 x 300 GB 15 k rpm 3.5-inch SAS disks For 1,000 virtual desktops • 5 x 300 GB 15 k rpm 3.5-inch SAS disks For 2,000 virtual desktops • 10 x 300 GB 15 k rpm 3.5-inch SAS disks	Optional for vCenter Operations Manager for View

Note: EMC recommends configuring at least one hot spare for every 30 drives of a given type. The recommendations above do not include hot spares.

vSphere Storage Virtualization for VSPEX

This section provides guidelines for setting up the storage layer of the solution to provide high availability and the expected level of performance.

VMware ESXi provides host-level storage virtualization. It virtualizes the physical storage and presents the virtualized storage to the virtual machine.

A virtual machine stores its operating system and all other files related to the virtual machine activities in a virtual disk. The virtual disk can be one file or multiple files. VMware uses a virtual SCSI controller to present the virtual disk to the guest operating system running inside the virtual machine. Figure 22 on page 83 shows various VMware virtual disk types.

The virtual disk resides in a datastore. Depending on the type used, it can be either a VMware Virtual Machine File system (VMFS) datastore or an NFS datastore.

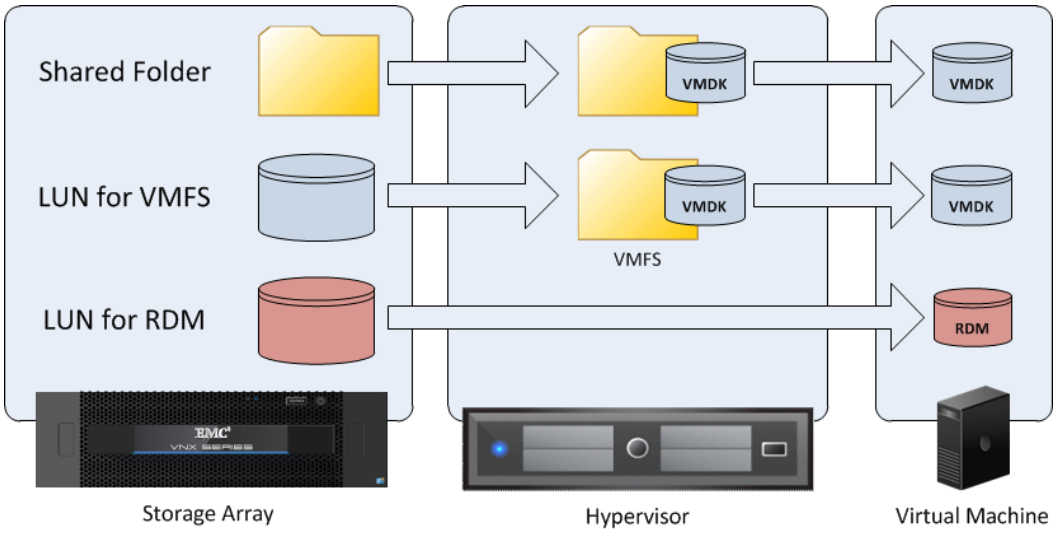

Figure 22. VMware virtual disk types

VMFS

VMFS is a cluster file system that provides storage virtualization optimized for virtual machines. It can be deployed over any SCSI-based local or network storage.

Raw device mapping

VMware also provides a mechanism called raw device mapping (RDM), which uses a Fibre Channel or iSCSI protocol and allows a virtual machine to have direct access to a volume on the physical storage.

NFS

VMware supports using NFS file systems from an external NAS storage system or device as a virtual machine datastore.

VSPEX storage building block

Sizing the storage system to meet virtual server IOPS is a complicated process. When an I/O reaches the storage array, several components, such as the Data Mover (for file-based storage), SPs, back-end dynamic random access memory (DRAM) cache, FAST Cache (if used), and disks, serve that I/O. Customers must consider various factors when planning and scaling their storage system to balance capacity, performance, and cost for their applications.

VSPEX uses a building block approach to reduce complexity. A building block is a set of disk spindles that can support a certain number of virtual desktops in the VSPEX architecture. Each building block combines several disk spindles to create a storage pool that supports the needs of the private cloud environment. Each building block storage pool, regardless of the size, contains two flash drives with FAST VP storage tiering to enhance metadata operations and performance.

Building block for 500 virtual desktops

The first building block can contains up to 500 virtual desktops with ten SAS drives in a FAST Cache enabled storage pool, as shown in Figure 23.

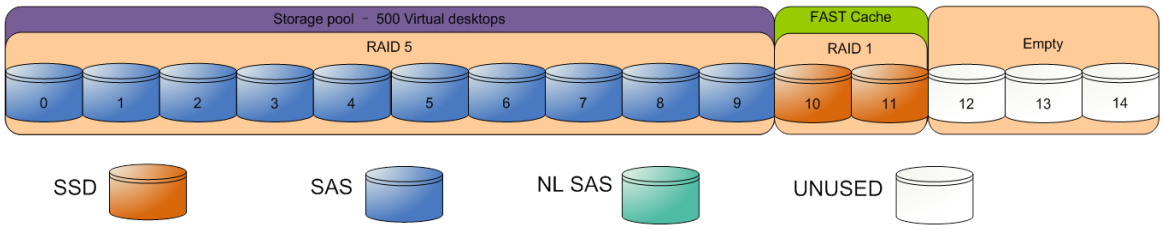

Figure 23. Storage layout building block for 500 virtual desktops

This is the smallest building block qualified for the VSPEX architecture.

Building block for 1,000 virtual desktops

The second building block can contain up to 1,000 virtual desktops. It contains 15 SAS drives in a FAST Cache enabled pool as shown in Figure 24.

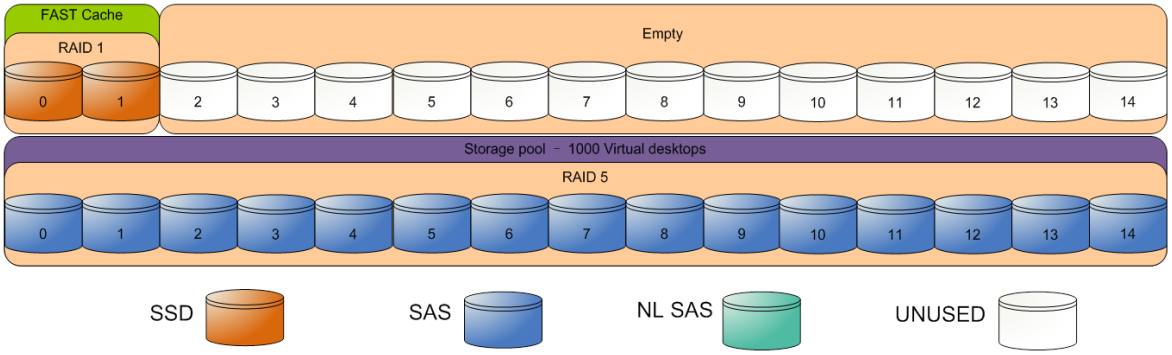

Figure 24. Storage layout building block for 1,000 virtual desktops

These two building blocks are currently verified on the VNX series and provide a flexible solution for VSPEX sizing. Table 12 shows a simple list of the disks required to support different configuration scales, excluding hot spare needs.

Note: If a configuration is started with the 500 desktop building block, it can be expanded to the 1,000 desktop building block by adding five matching SAS drives and allowing the pool to restripe.

Table 12. Number of disks required for different number of virtual desktops

Virtual desktops	Flash drives (FAST Cache)	SAS drives
500	2	10
1,000	2	15
2,000	4	30

VSPEX end user computing validated maximums

VSPEX end user computing configurations are validated on the VNX5400 and VNX5600 platforms. Each platform has different capabilities in terms of processors, memory, and disks. For each array, there is a recommended maximum VSPEX end-user-computing configuration, which is demonstrated below. Smaller implementations are supported using the building block method described in the previous section.

VNX5400

Core storage layout

Figure 25 illustrates the layout of the disks that are required to store 1,000 desktop virtual machines. This layout does not include space for user profile data. Refer to VNX shared file systems for more information.

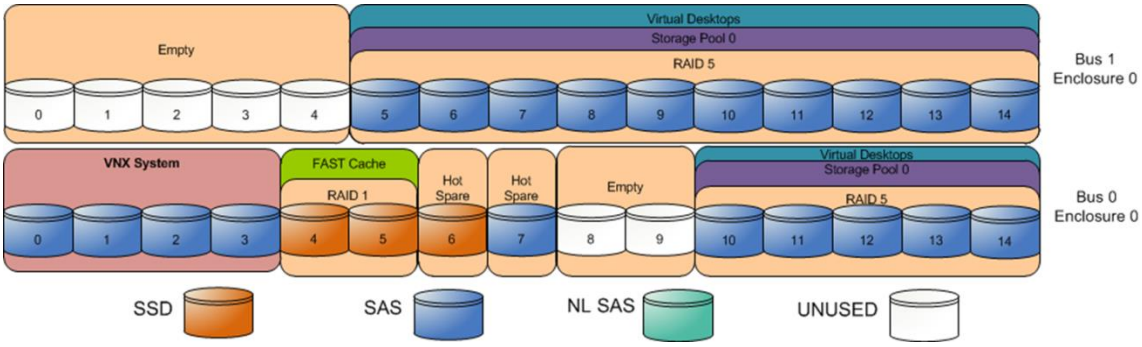

Figure 25. Core storage layout for 1,000 virtual desktops using VNX5400

Core storage layout overview

The following core configuration is used in the solution:

- Four SAS disks (0_0_0 to 0_0_3) are used for the VNX OE.
- Disks shown as 0_0_4 and 1_0_0 are hot spares.
- Fifteen SAS disks (shown as 0_0_10 to 0_0_14 and 1_0_5 to 1_0_14) in the RAID 5 storage pool 0 are used to store virtual desktops. FAST Cache is enabled for the entire pool.

▪ For file storage, 10 LUNs of 305 GB each are provisioned from the pool to provide the storage required to create eight 368 GB NFS file systems and two 50 GB file systems. The file systems are presented to the vSphere servers as NFS datastores.

- For block storage, 8 LUNs of 369 GB each and 2 LUNs of 50 GB each are provisioned from the pool to present to the vSphere servers as 10 VMFS datastores.

Note: Two 50 GB datastores are used to save replica disks.

- Two flash drives (shown as 0_0_4 to 0_0_5) are used for EMC VNX FAST Cache. There are no user-configurable LUNs on these drives.
- Disks shown as 0_0_8 to 0_0_9 and 1_0_0 to 1_0_4 are unused and were not used for testing this solution.

Note: If more capacity is desired, larger drives may be substituted. To satisfy the load recommendations, the drives will all need to be 15 k rpm and the same size. If you use different sizes, storage layout algorithms may give sub-optimal results.

Optional user data storage layout

In solution validation testing, storage space for user data was allocated on the VNX array as shown in Figure 26. This storage is in addition to the core storage shown above. If storage for user data exists elsewhere in the production environment, this storage is not required.

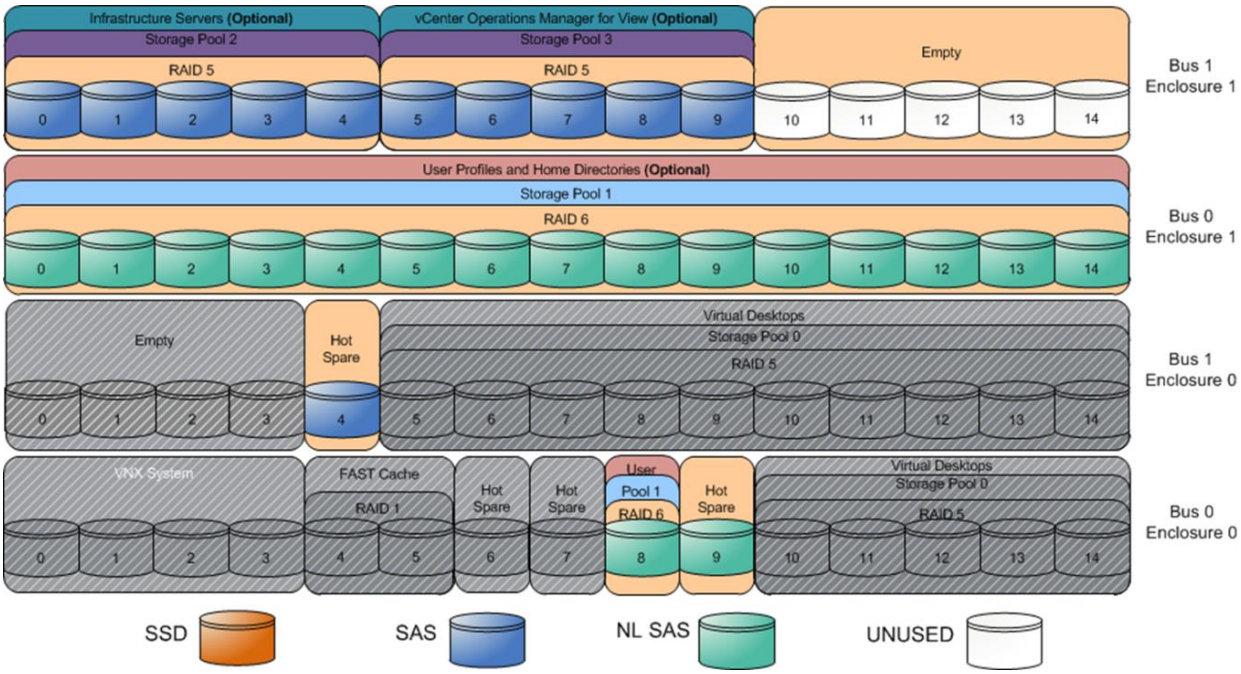

Figure 26. Optional storage layout for 1,000 virtual desktops using VNX5400

Optional storage layout overview

The following optional configuration is used in the solution:

- Disks shown as 0_0_9 and 1_0_4 are hot spares. These disks are marked as hot spare in the storage layout diagram.

- Five SAS disks (shown as 1_1_0 to 1_1_4) in the RAID 5 Storage Pool 2 are used to store the infrastructure virtual machines. A 1.0 TB LUN or NFS file system is provisioned from the pool to present to the vSphere servers as a VMFS or NFS datastore.
- Five SAS disks (shown as 1_1_5 to 1_1_9) in the RAID 5 Storage Pool 3 are used to store the vCenter Operations Manager for View virtual machines and databases. A 1.0 TB LUN or NFS file system is provisioned from the pool to present to the vSphere servers as a VMFS or NFS datastore.
- Sixteen NL-SAS disks (shown as 0_0_8 and 0_1_0 to 0_1_14) in the RAID 6 Storage Pool 1 are used to store user data and profiles. Ten LUNs of 600 GB each are provisioned from the pool to provide the storage required to create four CIFS file systems.
- Disks shown as 1_0_0 to 1_0_3 and 1_1_10 to 1_1_14 are unused and not used for testing this solution.
- Disks shaded gray are required and part of the core storage layout.

If multiple drive types have been implemented, FAST VP may be enabled to automatically tier data to balance differences in performance and capacity.

Note: Do not use FAST VP for virtual desktop datastores. FAST VP can provide performance improvements when implemented for user data and roaming profiles.

VNX shared file systems

Virtual desktops use four shared file systems—two for the VMware Horizon View Persona Management repositories and two to redirect user storage that resides in home directories. In general, redirecting users' data out of the base image to VNX for File enables centralized administration, backup and recovery, and makes the desktops more stateless. Each file system is exported to the environment through a CIFS share. Each persona management repository share and home directory share serves 500 users.

VNX5600

Core storage layout

Figure 27 illustrates the layout of the disks that are required to store 2,000 desktop virtual machines. This layout does not include space for user profile data. Refer to VNX shared file systems for more information.

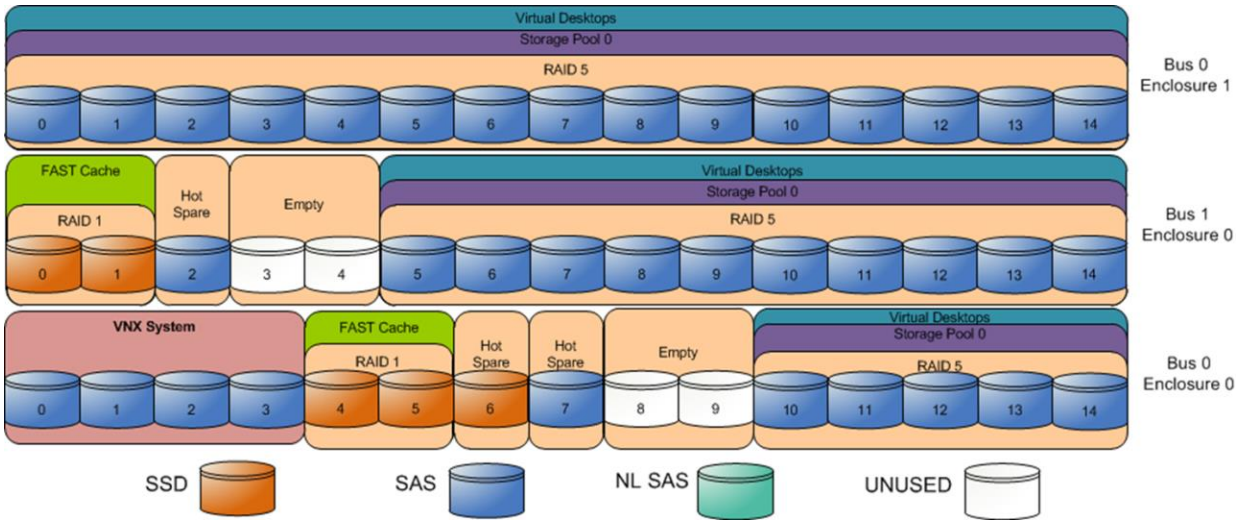

Figure 27. Core storage layout for 2,000 virtual desktops using VNX5600

Core storage layout overview

The following core configuration is used in the solution:

- Four SAS disks (shown as 0_0_0 to 0_0_3) are used for the VNX OE.
- Disks shown as 0_0_6 and 0_0_7 are hot spares. These disks are marked as hot spare in the storage layout diagram.
- Fifteen SAS disks (shown as 0_0_10 to 0_0_14 and 1_0_5 to 1_0_14) in the RAID 5 Storage Pool 0 are used to store virtual desktops. FAST Cache is enabled for the entire pool.

- For file storage, 10 LUNs of 611 GB each are provisioned from the pool to provide the storage required to create sixteen 375 GB NFS file systems and two 50 GB file systems. The file systems are presented to the vSphere servers as NFS datastores.
- For block storage, 16 LUNs of 375 GB each and 2 LUNs of 50 GB each are provisioned from the pool to present to the vSphere servers as 10 VMFS datastores.

Note: Two 50 GB datastores are used to save replica disks

- Four flash drives (shown as 0_0_4 to 0_0_5 and 1_0_0 to 1_0_1) are used for EMC VNX FAST Cache. There are no user-configurable LUNs on these drives.
- Disks shown as 0_0_8 to 0_0_9 and 1_0_3 to 1_0_4 were not used for testing this solution.

Note: If more capacity is desired, larger drives may be substituted. To satisfy the load recommendations, the drives will all need to be 15 k rpm and the same size. If differing sizes are used, storage layout algorithms may give sub-optimal results.

Optional user data storage layout

In solution validation testing, storage space for user data was allocated on the VNX array as shown in Figure 28. This storage is in addition to the core storage shown above. If storage for user data exists elsewhere in the production environment, this storage is not required.

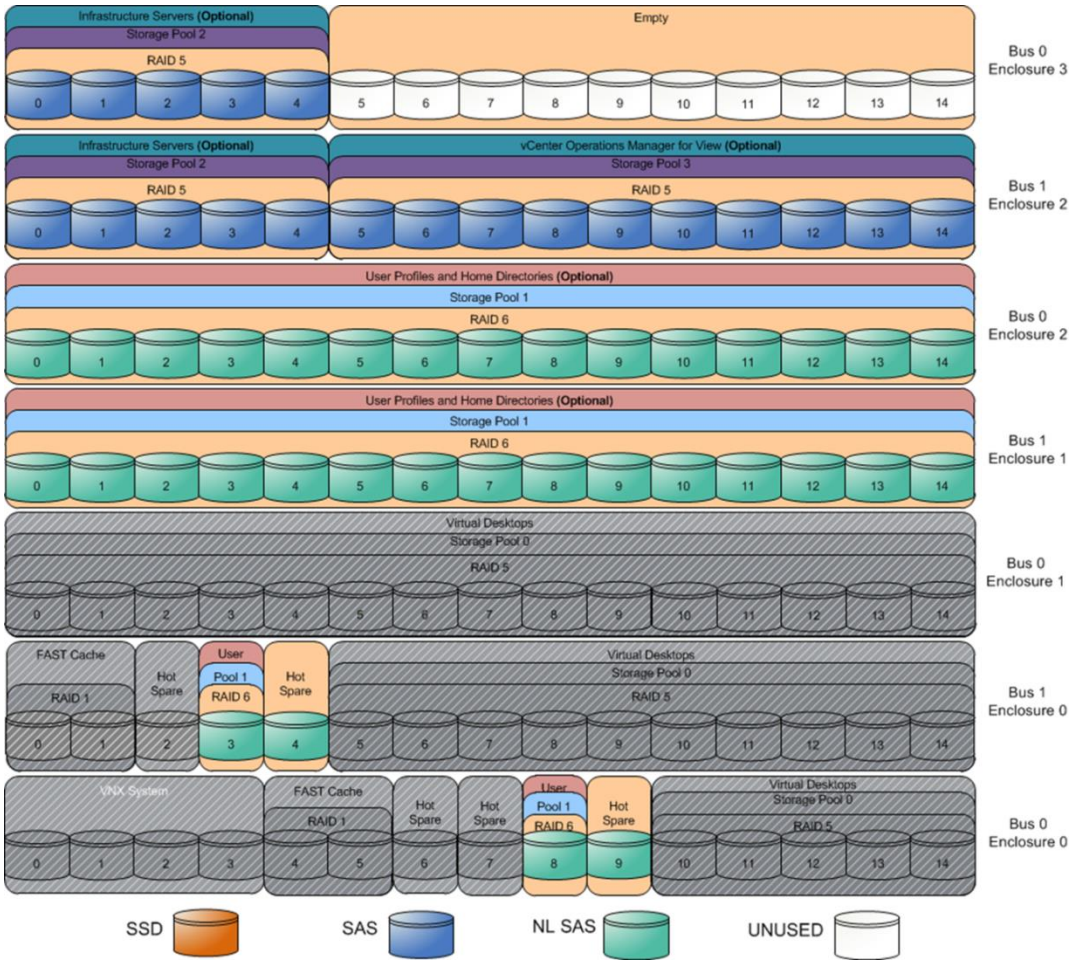

Figure 28. Optional storage layout for 2,000 virtual desktops using VNX5600

Optional storage layout overview

The following optional configuration is used in the solution:

- Disks shown as 0_0_9 and 1_0_4 are hot spares. These disks are marked as hot spare in the storage layout diagram.
- Five SAS disks (shown as 1_2_0 to 1_2_4) in the RAID 5 Storage Pool 2 are used to store the infrastructure virtual machines. A 1.0 TB LUN or NFS file system is provisioned from the pool to present to the vSphere servers as a VMFS or NFS datastore.
- Ten SAS disks (shown as 1_2_5 to 1_2_14) in the RAID 5 Storage Pool 3 are used to store the vCenter Operations Manager for View virtual machines and databases. A 2.0 TB LUN or NFS file

system is provisioned from the pool to present to the vSphere servers as a VMFS or NFS datastore.

- Thirty-two NL-SAS disks (shown as 0_0_8, 1_0_3, 1_1_0 to 1_1_14, and 0_2_0 to 0_2_14) in the RAID 6 Storage Pool 1 are used to store user data and profiles. FAST Cache is enabled for the entire pool. Ten LUNs of 3 TB each are provisioned from the pool to provide the storage required to create four CIFS file systems.
- Disks shown as 1_2_10 to 1_2_14 are unbound and not used for testing this solution.
- Disks shaded gray are required and are part of the core storage layout.

If multiple drive types have been implemented, FAST VP may be enabled to automatically tier data to balance differences in performance and capacity.

VNX shared file systems

Four shared file systems are used by the virtual desktops—two for the VMware View Persona Management repositories and two to redirect user storage that resides in home directories. In general, redirecting users' data out of the base image to VNX for File enables centralized administration, backup and recovery, and makes the desktops more stateless. Each file system is exported to the environment through a CIFS share. Each persona management repository share and home directory share serves 1,000 users.

High availability and failover

Introduction

This VSPEX solution provides a highly available virtualized server, network, and storage infrastructure. When implemented in accordance with this guide it provides the ability to survive single-unit failures with minimal impact to business operations.

Virtualization layer

EMC recommends configuring high availability in the virtualization layer and automatically allowing the hypervisor to restart virtual machines that fail. Figure 29 illustrates the hypervisor layer responding to a failure in the compute layer.

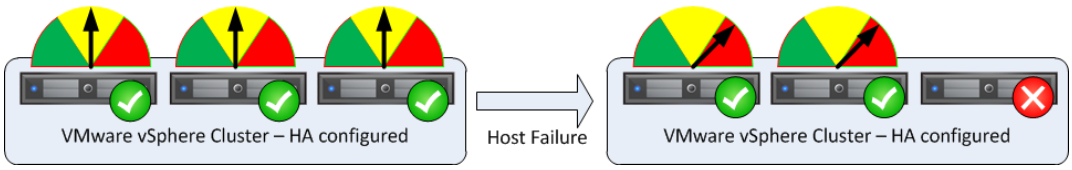

Figure 29. High availability at the virtualization layer

By implementing high availability at the virtualization layer, even in the event of a hardware failure, the infrastructure will attempt to keep as many services running as possible.

Compute layer

While the choice of servers to implement in the compute layer is flexible, it is best to use the enterprise class servers designed for data centers. This type of server has redundant power supplies, as shown in Figure 30. You should connect them to separate Power Distribution Units (PDUs) in accordance with your server vendor's best practices.

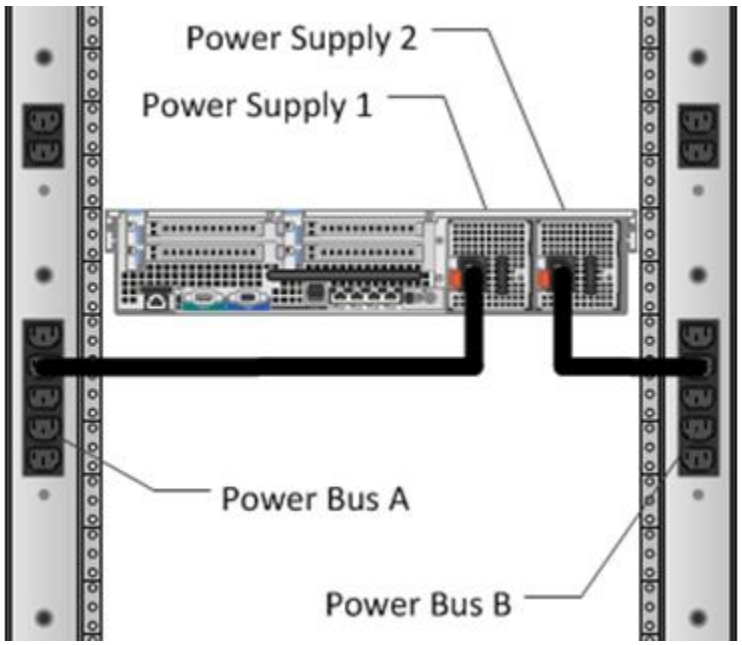

Figure 30. Redundant power supplies

We also recommend that you configure high availability in the virtualization layer. This means that you must configure the compute layer with enough resources to ensure that the total number of available resources meets the needs of the environment, even with a server failure. Figure 29 demonstrates this recommendation.

Brocade Network layer

The advanced networking features of the VNX storage family and Brocade VDX Ethernet Fabric or Connectrix-B Fibre Channel Family of switches provide protection against network connection failures at the array. Each vSphere host has multiple connections to user and storage Ethernet networks to guard against link failures. These connections should be spread across multiple Ethernet switches to guard against component failure in the network. These connections are illustrated in Figure 31 and Figure 32.

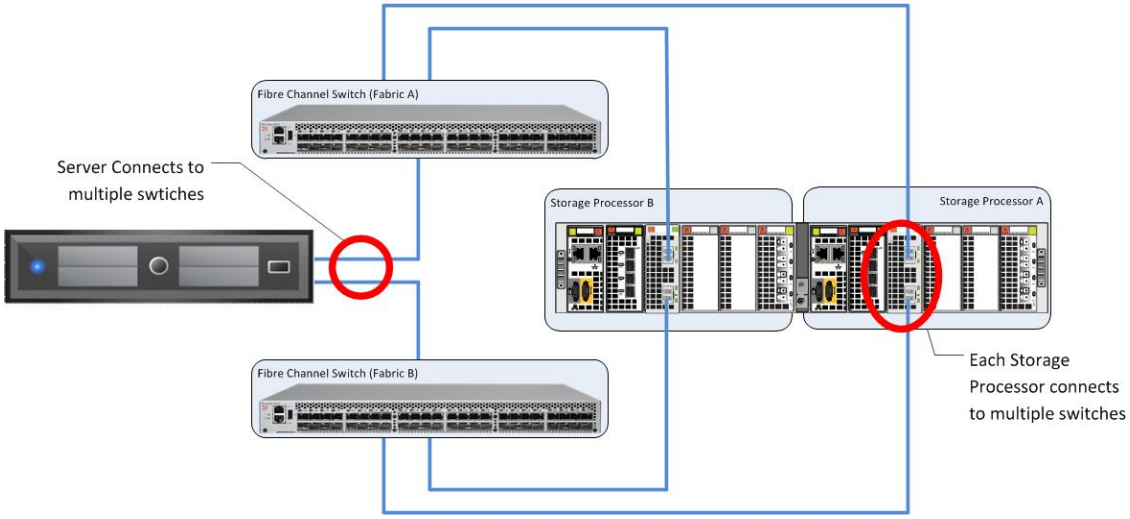

Figure 31. Brocade Network layer High-Availability (VNX) – block storage network variant

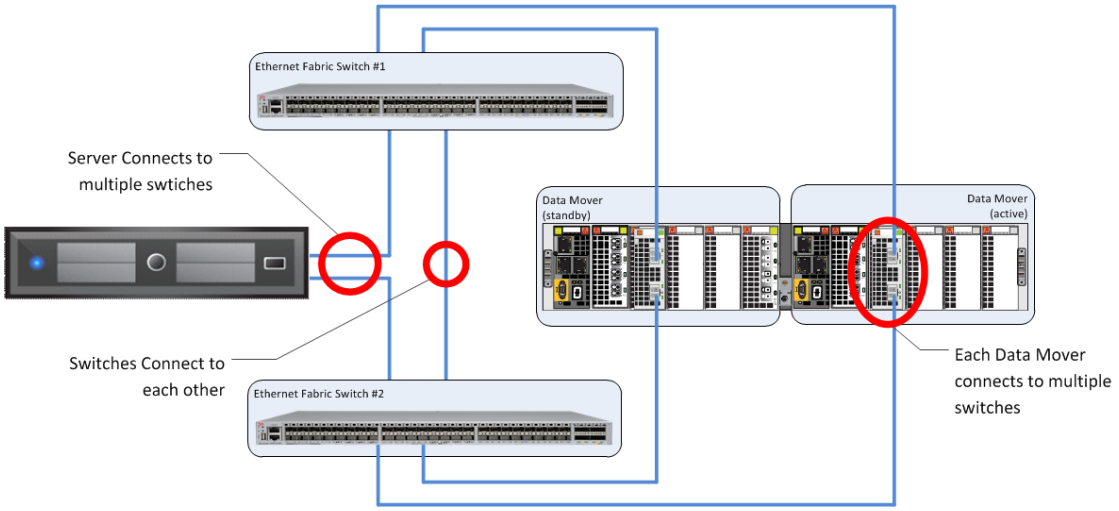

Figure 32. Brocade Network layer High-Availability (VNX) - file storage network variant

By ensuring that there are no single points of failure in the network layer you can ensure that the compute layer is able to access storage and communicate with users even if a component fails.

Storage layer

The VNX family is designed for five 9s availability by using redundant components throughout the array as shown in Figure 33. All of the array components are capable of continued operation in case of hardware failure. The RAID disk configuration on the array provides protection against data loss due to individual disk failures, and you can dynamically allocate the available hot spare drives to replace a failing disk.

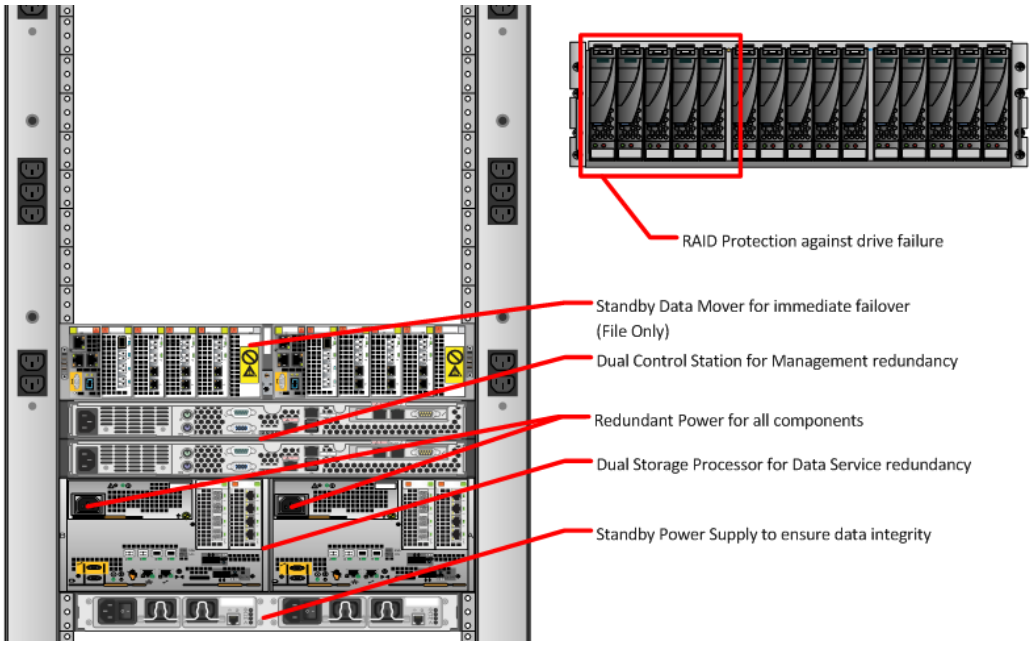
RAID Protection against drive failure
Standby Data Mover for immediate failover
(File Only)
Dual Control Station for Management redundancy
Redundant Power for all components
Dual Storage Processor for Data Service redundancy
Standby Power Supply to ensure data integrity

Figure 33. VNX series high availability

EMC storage arrays are designed to be highly available by default. Use the installation guides to ensure that there are no single unit failures that result in data loss or unavailability.

Validation test profile

Profile characteristics

Table 13 shows the solution stacks that we validated with the environment profile.

Table 13. Validated environment profile

Profile characteristic	Value
Virtual desktop operating system	Windows 7 Enterprise (32-bit) SP1
CPU per virtual desktop	1 vCPU
Number of virtual desktops per CPU core	8
RAM per virtual desktop	2 GB
Desktop provisioning method	Linked Clones
Average storage available for each virtual desktop	3 GB (vmdk and vswap)
Average IOPS per virtual desktop at steady state	10 IOPS
Average peak IOPS per virtual desktop during boot storm	14 IOPS (File storage) 23 IOPS (Block storage)
Number of datastores to store virtual desktops	4 for 500 virtual desktops 8 for 1,000 virtual desktops 16 for 2,000 virtual desktops
Number of virtual desktops per datastore	125
Disk and RAID type for datastores	RAID 5, 300 GB, 15 k rpm, 3.5-inch SAS disks
Disk and RAID type for CIFS shares to host user profiles and home directories (optional)	RAID 6, 2 TB, 7,200 rpm, 3.5-inch NL-SAS disks

Antivirus and antimalware platform profile

Platform characteristics

Table 14 shows how the solution was sized based on the following vShield Endpoint platform requirements.

Table 14. Platform characteristics

Platform Component	Technical Information
VMware vShield Manager appliance	Manages the vShield Endpoint service installed on each vSphere host
	1 vCPU, 3 GB RAM, and 8 GB hard disk space
VMware vShield Endpoint service	Installed on each desktop vSphere host. The service uses up to 512 MB of RAM on the vSphere host.
VMware Tools vShield Endpoint component	A component of the VMware tools suite that enables integration with the vSphere host vShield Endpoint service
	The vShield Endpoint component of VMware tools is installed as an optional component of the VMware tools software package and should be installed on the master virtual desktop image.
vShield Endpoint third-party security plug-in	Requirements vary based on individual vendor specifications. Refer to the selected third party vendor documentation to understand what resources are required. **Note:** A third party plug-in and associated components are required to complete the vShield Endpoint solution. Refer to vendor documentation for specific details concerning vShield Endpoint requirements.

vShield Architecture

The individual components of the VMware vShield Endpoint platform and the vShield partner security plug-ins each have specific CPU, RAM, and disk space requirements. The resource requirements vary based on a number of factors such as the number of events being logged, log retention needs, the number of desktops being monitored, and the number of desktops present on each vSphere host.

vCenter Operations Manager for View platform profile desktops

Platform characteristics

Table 15 shows how this solution stack was sized based on the following vCenter Operations Manager for View platform requirements.

Table 15. Platform characteristics

Platform Component	Technical Information
VMware vCenter Operations Manager vApp	The vApp consists of a user interface (UI) virtual appliance and an Analytics virtual appliance. For 500 virtual desktops • **UI appliance requirements:** 2 vCPU, 5 GB RAM, and 50 GB hard disk space • **Analytics appliance requirements:** 2 vCPU, 7 GB RAM, and 300 GB hard disk space For 1,000 virtual desktops • **UI appliance requirements:** 2 vCPU, 7 GB RAM, and 75 GB hard disk space. • **Analytics appliance requirements:** 2 vCPU, 9 GB RAM, and 600 GB hard disk space. For 2,000 virtual desktops • **UI appliance requirements:** 4 vCPU, 11 GB RAM, and 150 GB hard disk space. • **Analytics appliance requirements:** 4 vCPU, 14 GB RAM, and 1.2 TB hard disk space.
VMware vCOps for View Adapter	The vCOps for View Adapter enables integration between vCenter Operations Manager and VMware Horizon View and requires a server running Microsoft Windows 2008 R2. The adapter gathers View related status information and statistical data. For 500 virtual desktops • **Server requirements**: 2 vCPU, 6 GB RAM, and 30 GB hard disk space. For 1,000 virtual desktops • **Server requirements:** 2 vCPU, 6 GB RAM, and 30 GB hard disk space. For 2,000 virtual desktops • **Server requirements**: 4 vCPU, 8 GB RAM, and 30 GB hard disk space

Platform Component	Technical Information
VMware vCenter Operations Manager vApp	The vApp consists of a user interface (UI) virtual appliance and an Analytics virtual appliance. For 500 virtual desktops • **UI appliance requirements**: 2 vCPU, 5 GB RAM, and 50 GB hard disk space • **Analytics appliance requirements**: 2 vCPU, 7 GB RAM, and 300 GB hard disk space For 1,000 virtual desktops • **UI appliance requirements:** 2 vCPU, 7 GB RAM, and 75 GB hard disk space. • **Analytics appliance requirements:** 2 vCPU, 9 GB RAM, and 600 GB hard disk space. For 2,000 virtual desktops • **UI appliance requirements:** 4 vCPU, 11 GB RAM, and 150 GB hard disk space. • **Analytics appliance requirements:** 4 vCPU, 14 GB RAM, and 1.2 TB hard disk space.
VMware vCOps for View Adapter	The vCOps for View Adapter enables integration between vCenter Operations Manager and VMware Horizon View and requires a server running Microsoft Windows 2008 R2. The adapter gathers View related status information and statistical data. For 500 virtual desktops • **Server requirements**: 2 vCPU, 6 GB RAM, and 30 GB hard disk space. For 1,000 virtual desktops • **Server requirements:** 2 vCPU, 6 GB RAM, and 30 GB hard disk space. For 2,000 virtual desktops • **Server requirements**: 4 vCPU, 8 GB RAM, and 30 GB hard disk space

vCenter Operations Manager for View Architecture

The individual components of vCenter Operations Manager for View have specific CPU, RAM, and disk space requirements. The resource requirements vary based on the number of desktops being monitored. The numbers provided in Table 15 assume that 500, 1,000, and 2,000 desktops will be monitored.

Backup and recovery configuration guidelines

See Design and Implementation Guide: EMC Backup and Recovery Options for VSPEX End User Computing for VMware Horizon View on available on EMC Online Support.

Sizing guidelines

The following sections define the reference workload used to size and implement the VSPEX architectures discussed in this document, and provide guidance on how to correlate those reference workloads to actual customer workloads. They also explain how that may change the end delivery from the server and network perspective.

You can modify the storage definition by adding drives for greater capacity and performance and adding features like FAST Cache for desktops and FAST VP for improved user data performance. The disk layouts have been created to provide support for the appropriate number of virtual desktops at the defined performance level. Decreasing the number of recommended drives or stepping down an array type can result in lower IOPS per desktop and a less satisfactory user experience due to higher response times.

Reference workload

Each VSPEX Proven Infrastructure balances the storage, network, and compute resources needed for a set number of virtual machines that have been validated by EMC. In practice, each virtual machine has its own set of requirements that rarely fit a pre-defined idea of what a virtual machine should be. In any discussion about virtual infrastructures, define a reference workload first. Not all servers perform the same tasks, and it is impractical to build a reference that takes into account every possible combination of workload characteristics.

Defining the reference workload

To simplify the discussion, we have defined a representative customer reference workload. By comparing your actual customer usage to this reference workload, you can extrapolate which reference architecture to choose.

For the VSPEX end-user computing solutions, the reference workload is defined as a single virtual desktop. Table 16 shows the reference virtual machine has following detail characteristics:

Table 16. Virtual desktop characteristics

Characteristic	Value
Virtual desktop operating system	Microsoft Windows 7 Enterprise Edition (32-bit) SP1
Virtual processors per virtual desktop	1
RAM per virtual desktop	2 GB
Available storage capacity per virtual desktop*	3 GB (vmdk and vswap)
Average IOPS per virtual desktop at steady state	10
* This available storage capacity is calculated based on drives used in this solution. You can create more space by adding drives or using larger capacity drives of the same class.	

This desktop definition is based on user data that resides on shared storage. The I/O profile is defined by using a test framework that runs all desktops concurrently with a steady load generated by the constant use of office-based applications like browsers, office productivity software, and other standard task worker utilities.

Applying the reference workload

In addition to the supported desktop numbers (1,000, 2,000, or 3,000), there may be other factors to consider when deciding which end-user computing solution to deploy.

Concurrency

The workloads used to validate VSPEX solutions assume that all desktop users will be active at all times. In other words, the 1,000-desktop architecture was tested with 1,000 desktops, all generating workload in parallel, all booted at the same time, and so on. If the customer expects to have 1,200 users, but only 50 percent of them will be logged on at any given time due to time zone differences or alternate shifts, the 600 active users out of the total 1,200 users can be supported by the 1,000-desktop architecture.

Heavier desktop workloads

The workload defined in Table 16 and used to test these VSPEX End-User Computing solution configurations is considered a typical office worker load. However, some customers' users might have a more active profile.

If a company has 800 users and, due to custom corporate applications, each user generates 15 IOPS as compared to 10 IOPS used in the VSPEX workload, this customer will need 12,000 IOPS (800 users * 15 IOPS per desktop). The 1,000-desktop configuration would be underpowered in this case because it has been rated to 10,000 IOPS (1,000 desktops * 10 IOPS per desktop). This customer should consider moving up to the 2,000-desktop solution.

Implementing the reference architectures

Overview

The solutions architectures require a set of hardware to be available for the CPU, memory, network, and storage needs of the system. In the solutions architectures, these are presented as general requirements that are independent of any particular implementation. This section describes some considerations for implementing the requirements.

Resource types

The solution architectures define the hardware requirements for the solution in terms of five basic types of resources:

- CPU resources
- Memory resources
- Network Resources
- Storage resources
- Backup resources

This section describes the resource types, how they are used in the solution, and key considerations for implementing them in a customer environment.

CPU resources

The architectures define the number of CPU cores that are required, but not a specific type or configuration. It is intended that new deployments use recent revisions of common processor technologies. It is assumed that these will perform as well as, or better than, the systems used to validate the solution.

When using Avamar backup solution for VSPEX, do not schedule all backups at once; instead, stagger them across your backup window. Scheduling all resources to back up at the same time could cause the consumption of all available host CPU.

In any running system, monitor the utilization of resources and adapt as needed. The reference virtual desktop and required hardware resources in the solutions assume that there will be no more than eight virtual CPUs for each physical processor core (8:1 ratio). In most cases, this provides an appropriate level of resources for the hosted virtual desktops; however, this ratio may not be appropriate in all use cases. EMC recommends monitoring the CPU utilization at the hypervisor layer to determine if more resources are required.

Memory resources

Each virtual desktop in the solution is defined to have 2 GB of memory. In a virtual environment, because of budget constraints, it is common to provision virtual desktops with more memory than the hypervisor physically has. The memory over-commitment technique takes advantage of the fact that each virtual desktop does not fully utilize the amount of memory allocated to it. It makes business sense to oversubscribe the memory usage

to some degree. The administrator has the responsibility to monitor proactively the oversubscription rate such that it does not shift the bottleneck away from the server and become a burden to the storage subsystem.

If VMware vSphere runs out of memory for the guest operating systems, paging will begin to take place, resulting in extra I/O activity going to the vswap files. If the storage subsystem is sized correctly, occasional spikes due to vswap activity might not cause performance issues because transient bursts of load can be absorbed. However, if the memory oversubscription rate is so high that the storage subsystem is severely impacted by a continuing overload of vswap activity, more disks will need to be added—not because of capacity requirement—but due to the demand of increased performance. The administrator must decide whether it is more cost effective to add more physical memory to the server or to increase the amount of storage. With memory modules being a commodity, it is likely less expensive to choose the former option.

This solution was validated with statically assigned memory and no over-commitment of memory resources. If memory over-commit is used in a real-world environment, regularly monitor the system memory utilization and associated page file I/O activity to ensure that a memory shortfall does not cause unexpected results.

Network resources

The solution outlines the minimum needs of the system. If additional bandwidth is needed, it is important to add capability at both the storage array and the hypervisor host to meet the requirements. The options for network connectivity on the server will depend on the type of server. The storage arrays have a number of included network ports and have the option to add ports using EMC FLEX I/O modules.

For reference purposes in the validated environment, EMC assumes that each virtual desktop generates 10 I/Os per second with an average size of 4 KB. This means that each virtual desktop is generating at least 40 KB/s of traffic on the storage network. For an environment rated for 500 virtual desktops, this comes out to a minimum of approximately 20 MB/sec. This is well within the bounds of modern networks. However, this does not consider other operations. For example, additional bandwidth is needed for:

- User network traffic
- Virtual desktop migration
- Administrative and management operations

The requirements for each of these will vary depending on how the environment is being used. It is not practical to provide concrete numbers in this context. The network described in the solution for each solution should be sufficient to handle average workloads for the above use cases. The specific Brocade storage network layer connectivity solution is defined in Chapter 5.

Regardless of the network traffic requirements, always have at least two physical network connections that are shared for a logical network to ensure a single link failure does not affect the availability of the system. The network should be designed so that the aggregate bandwidth, in the event of a failure, is sufficient to accommodate the full workload.

Storage resources

The solutions contain layouts for the disks used in the validation of the system. Each layout balances the available storage capacity with the performance capability of the drives. There are a few layers to consider when examining storage sizing. Specifically, the array has a collection of disks that are assigned to a storage pool. From that storage pool, you can provision datastores to the VMware vSphere Cluster. Each layer has a specific configuration that is defined for the solution and documented in Chapter 5.

It is generally acceptable to replace drive types with a type that has more capacity with the same performance characteristic or with types that have higher performance characteristics and the same capacity. It is also acceptable to change the placement of drives in the drive shelves in order to comply with updated or new drive shelf arrangements.

In other cases, where there is a need to deviate from the proposed number and type of drives specified or the specified pool and datastore layouts, ensure that the target layout delivers the same or greater resources to the system.

Backup resources

See Design and Implementation Guide: EMC Backup and Recovery Options for VSPEX End User Computing for VMware Horizon View available on EMC Online Support.

Expanding existing VSPEX End User Ccomputing environments

The EMC VSPEX End User Computing solution supports a flexible implementation model where it is easy to expand your environment as the needs of the business change.

You can combine the building block configurations presented in this solution to form larger implementations. For example, the 1,000-desktop configuration can be realized by starting with that configuration, or by starting with the 500-desktop configuration and expanding it when needed. In the same way the 2,000-desktop configuration can be implemented all at once or gradually by expanding the storage resources as you need them.

Implementation summary

The requirements stated in the solution are what EMC considers the minimum set of resources to handle the workloads required based on the stated definition of a reference virtual desktop. In any customer implementation, the load of a system will vary over time as users interact with the system. However, if the customer virtual desktops differ significantly from the reference definition and vary in the same resource group, you may need to add more of that resource to the system.

Quick assessment

Overview

An assessment of the customer environment helps to ensure that you implement the correct VSPEX solution. This section provides an easy-to-use worksheet to simplify the sizing calculations, and help assess the customer environment.

First, summarize the user types planned for migration into the VSPEX End-User Computing environment. For each group, determine the number of virtual CPUs, the amount of memory, the required storage performance, the required storage capacity, and the number of reference virtual desktops required from the resource pool.

Applying the reference workload provides examples of this process.

Fill out a row in the worksheet for each application, as shown in Table 17.

Table 17. Blank worksheet row

Application		CPU (Virtual CPUs)	Memory (GB)	IOPS	Equivalent reference virtual desktops	Number of users	Total reference desktops
Example User Type	Resource Requirements						
	Equivalent Reference Desktops						

Fill out the resource requirements for the User Type. The row requires inputs on three different resources: CPU, Memory, and IOPS.

CPU requirements

The reference virtual desktop assumes most desktop applications are optimized for a single CPU. If one type of user requires a desktop with multiple virtual CPUs, modify the proposed virtual desktop count to account for the additional resources. For example, if you virtualize 100 desktops, but 20 users require 2 CPUs instead of 1, consider that your pool needs to provide 120 virtual desktops of capability.

Memory requirements

Memory plays a key role in ensuring application functionality and performance. Each group of desktops will have different targets for the available memory that is considered acceptable. Like the CPU calculation, if a group of users requires additional memory resources, simply adjust the number of planned desktops to accommodate the additional resource requirements.

For example, if there are 200 desktops to be virtualized, but each one needs 4 GB of memory instead of the 2 GB that is provided in the reference virtual desktop, plan for 400 reference virtual desktops.

Storage performance requirements

The storage performance requirements for desktops are usually the least understood aspect of performance. The reference virtual desktop uses a workload generated by an industry-recognized tool to execute a wide variety of office productivity applications that should be representative of the majority of virtual desktop implementations.

Storage capacity requirements

The storage capacity requirement for a desktop can vary widely depending on the types of applications in use and specific customer policies. The virtual desktops presented in this solution rely on additional shared storage for user profile data and user documents. This requirement is covered as an optional component that can be met with the addition of specific storage hardware defined in the solution. It can also be covered with existing file shares in the environment.

Determining equivalent reference virtual desktops

With all of the resources defined, determine an appropriate value for the Equivalent Reference virtual desktops line by using the relationships in Table 18. Round all values up to the closest whole number.

Table 18. Reference virtual desktop resources

Resource	Value for reference virtual desktop	Relationship between requirements and equivalent reference virtual desktops
CPU	1	Equivalent reference virtual desktops = Resource requirements
Memory	2	Equivalent reference virtual desktops = (Resource requirements)/2
IOPS	10	Equivalent reference virtual desktops = (Resource requirements)/10

For example, a group of 100 users need the 2 virtual CPUs and 12 IOPS per desktop described earlier along with 8 GB of memory on the resource requirements line. For this example they need 2 reference desktops of CPU, 4 reference desktops of memory, and 2 reference desktops of IOPS based on the virtual desktop characteristics in Table 18. These figures go in the "Equivalent Reference Desktops" row as shown in Table 19. Use the maximum value in the row to fill in the "Equivalent Reference Virtual Desktops" column.

Multiply the number of equivalent reference virtual desktops by the number of users to arrive at the total resource needs for that type of user.

Table 19. Example worksheet row

User Type		CPU (Virtual CPUs)	Memory (GB)	IOPS	Equivalent reference virtual desktops	Number of users	Total reference desktops
Heavy users	Resource requirements	2	8	12			
	Equivalent reference virtual desktops	2	4	2	4	100	400

Once the worksheet is filled out for each user type that the customer wants to migrate into the virtual infrastructure, compute the total number of reference virtual desktops required in the pool by computing the sum of the Total column on the right side of the worksheet as shown in Table 20.

Table 20. Example applications

User Type		CPU (virtual CPUs)	Memory (GB)	IOPS	Equivalent reference virtual desktops	Number of users	Total reference desktops
Heavy users	Resource requirements	2	8	12			
	Equivalent reference virtual desktops	2	4	2	4	100	400
Moderate users	Resource requirements	2	4	8			
	Equivalent reference virtual desktops	2	2	1	2	100	200
Typical users	Resource requirements	1	2	8			
	Equivalent reference virtual desktops	1	1	1	1	300	300
Total							900

The VSPEX End-User Computing Solutions define discrete resource pool sizes. For this solution set, the pool contains 500, 1,000, or 2,000. In the case of Table 20, the customer requires 900 virtual desktops of capability from the pool. The 1,000 virtual desktop resource pool provides sufficient resources for the current needs and room for growth.

Fine tuning hardware resources

In most cases, the recommended hardware for servers and storage will be sized appropriately based on the process described. However, in some cases you may want to further customize the hardware resources available to the system. A complete description of system architecture is beyond the scope of this document but you can customize your solution further at this point.

Storage resources

In some applications, there is a need to separate some storage workloads from other workloads. The storage layouts in the VSPEX architectures put all of the virtual desktops in a single resource pool. In order to achieve

workload separation, purchase additional disk drives for each group that needs workload isolation and add them to a dedicated pool.

It is not appropriate to reduce the size of the main storage resource pool in order to support isolation or to reduce the capability of the pool without additional guidance beyond this paper. The storage layouts presented in the solutions are designed to balance many different factors in terms of high availability, performance, and data protection. Changing the components of the pool can have significant and difficult-to-predict impacts on other areas of the system.

Server resources

For the server resources in the VSPEX end-user computing solution, it is possible to customize the hardware resources more effectively. To do this, first total the resource requirements for the server components as shown in Table 21. Note the addition of the "Total CPU Resources" and "Total Memory Resources" columns at the right of the table.

Table 21. Server resource component totals

User types		CPU (virtual CPUs)	Memory (GB)	Number of users	Total CPU resources	Total memory resources
Heavy users	Resource requirements	2	8	100	200	800
Moderate users	Resource requirements	2	4	100	200	400
Typical users	Resource requirements	1	2	300	300	600
Total					700	1,800

In this example, the target architecture required 700 virtual CPUs and 1,800 GB of memory. With the stated assumptions of 8 desktops per physical processor core and no memory over-provisioning, this translates to 88 physical processor cores and 1,800 GB of memory. In contrast, the 1,000 virtual desktop resource pool as documented in the solution calls for 2,000 GB of memory and at least 125 physical processor cores. In this environment, the solution can be effectively implemented with fewer server resources.

Note: Keep high availability requirements in mind when customizing the resource pool hardware.

Table 22 is a blank worksheet.

Table 22. Blank customer worksheet

User type		CPU (virtual CPUs)	Memory (GB)	IOPS	Equivalent reference virtual desktops	Number of users	Total reference desktops
	Resource requirements						
	Equivalent reference virtual desktops						
	Resource requirements						
	Equivalent reference virtual desktops						
	Resource requirements						
	Equivalent reference virtual desktops						
	Resource requirements						
	Equivalent reference virtual desktops						
Total							

Chapter 5 VSPEX Configuration Guidelines

This chapter presents the following topics:

Configuration overview 110
Pre-deployment tasks 111
Customer configuration data 114
Prepare, connect, and configure Brocade network switches 114
Configure Brocade VDX 6740 Switch (File Storage) 119
Prepare and configure the storage array 152
Install and configure vSphere hosts 162
Install and configure SQL Server database 167
VMware vCenter Server Deployment 171
Set Up VMware View Connection Server 174
Set Up EMC Avamar 178
Set up VMware vShield Endpoint 179
Set Up VMware vCenter Operations Manager for View 181
Summary 183

Configuration overview

Deployment process

The deployment process is divided into the stages shown in Table 23. Upon completion of the deployment, the VSPEX infrastructure will be ready for integration with the existing customer network and server infrastructure.

Table 23 lists the main stages in the solution deployment process. The table also includes references to chapters that provide relevant procedures.

Table 23. Deployment process overview

Stage	Description	Reference
1	Verify prerequisites	Pre-deployment tasks
2	Obtain the deployment tools	Pre-deployment tasks
3	Gather customer configuration data	Customer configuration data
4	Rack and cable the components	Refer to the vendor documentation
5	Install and configure the Brocade network switches , connect to the management and customer network	Prepare, connect, and configure Brocade network switches
6	Install and configure the VNX	Prepare and configure the storage
7	Configure virtual machine datastores	Prepare and configure the storage
8	Install and configure the servers	Install and configure vSphere hosts
9	Set up SQL Server (used by VMware vCenter and VMware Horizon View)	Install and configure SQL Server database
10	Install and configure vCenter and virtual machine networking	VMware vCenter Server Deployment
11	Set up VMware View Connection Server	Set Up VMware View Connection Server
12	Set up EMC Avamar	Set Up EMC Avamar Set Up EMC Avamar
13	Set up VMware vShield Endpoint	Set up VMware vShield Endpoint
14	Set up VMware vCenter Operations Manager (vCOps) for View	Set Up VMware vCenter Operations Manager for View

Pre-deployment tasks

Overview

Pre-deployment tasks, as shown in Table 25, include procedures that do not directly relate to environment installation and configuration, but you will need the results from these tasks at the time of installation. Examples of pre-deployment tasks are collection of host names, IP addresses, VLAN IDs, license keys, installation media, and so on. You should perform these tasks before the customer visit to reduce the amount of time required on site.

Table 24. Tasks for pre-deployment

Task	Description	Reference
Gather documents	Gather the related documents listed in Appendix C. These documents to provide detail on setup procedures and deployment best practices for the various components of the solution.	EMC documentation Brocade Documentation
Gather tools	Gather the required and optional tools for the deployment. Use Table 25 to confirm that all equipment, software, and appropriate licenses are available before the deployment process.	Table 25
Gather data	Collect the customer-specific configuration data for networking, naming, and required accounts. Enter this information into the Customer Configuration Data worksheet for reference during the deployment process.	Appendix B

Deployment prerequisites

Complete the *VNX Block Configuration Worksheet* for FC variant or *VNX File and Unified Worksheet* for NFS variant, available on EMC Online Support, to provide the most comprehensive array-specific information.

Table 25 itemizes the hardware, software, and license requirements to configure the solution. Visit EMC Online Support for more information on these prerequisites.

Table 25. Deployment prerequisites checklist

Requirement	Description
Hardware	Physical servers to host virtual desktops for sufficient physical server capacity to host desktops
	VMware vSphere servers to host virtual infrastructure servers **Note:** This requirement may be covered by existing infrastructure.
	Networking: Switch port capacity and capabilities as required by the End-User Computing. Brocade VDX 6740 switches (File based storage network connectivity) or Brocade 6510 switches (Block based storage network connectivity)
	EMC VNX multiprotocol storage array with the required disk layout
Software	VMware vSphere installation media
	VMware vCenter Server 5.5 installation media
	VMware vShield Manager Open Virtualization Appliance (OVA) file
	VMware vCenter Operations Manager OVA file
	VMware vCOps for View Adapter
	VMware Horizon View 5.3 installation media
	vShield Endpoint partner anti-virus solution management server software
	vShield Endpoint partner security virtual machine software
	EMC VSI for VMware vSphere Unified Storage Management
	EMC VSI for VMware vSphere Storage Viewer
	Brocade VDX vCenter plug-in

Requirement	Description
	Microsoft Windows Server 2008 R2 installation media (suggested OS for VMware vCenter and VMware View Connection Server)
	Microsoft Windows 7 SP1 installation media
	Microsoft SQL Server 2008 or later installation media **Note:** This requirement might be covered in the existing infrastructure.
Software–FC variant only	EMC PowerPath® Viewer
	EMC PowerPath Virtual Edition
Software–NFS variant only	EMC vStorage API for Array Integration plug-in
Licenses	VMware vCenter 5.5 license key
	VMware vSphere Desktop license keys
	VMware Horizon View Premier 5.5 license keys
	VMware vShield Endpoint license keys (VMware)
	VMware vShield Endpoint license keys (vShield Partner)
	VMware vCOps Note: Brocade vCOps plug-in available for Brocade Connectrix-B Fibre Channel switches
	Microsoft Windows Server 2008 R2 Standard (or later) license keys **Note:** This requirement might be covered in the existing Microsoft Key Management Server (KMS).
	Microsoft Windows 7 license keys **Note:** This requirement might be covered in the existing Microsoft Key Management Server (KMS).
	Microsoft SQL Server license key **Note:** This requirement might be covered in the existing infrastructure.
	40GbE Port Upgrade Licenses for Brocade VDX 6740 switches with NOS v4.1.0
Licenses FC variant only	EMC PowerPath Virtual Edition license files

Customer configuration data

To reduce the onsite time, you should assemble information such as IP addresses and hostnames as part of the planning process.

Appendix B provides a table enabling you to maintain a record of relevant information. You can expand or shorten this form as required while adding, modifying, and recording your deployment progress.

Additionally, complete the *VNX Block Configuration Worksheet for FC variant* or *VNX File and Unified Worksheet for NFS variant,* available on EMC online support, to provide the most comprehensive array-specific information.

Prepare, connect, and configure Brocade network switches

Overview

This section lists the Brocade network infrastructure required to support the VSPEX architectures. Table 26 provides a summary of the tasks for switch and network configuration, and references for further information.

Table 26. Tasks for switch and network configuration

Task	Description	Reference
Complete network cabling	Connect switch interconnect ports. Connect VNX ports. Connect vSphere server ports.	
Configure Brocade Storage Network	Configure Brocade VDX Ethernet Fabric Network (File based Storage Network)	Configure Brocade VDX Ethernet Fabric Network (File based Storage Network)
	Configure Brocade 6510 Fabric Network (Block based Storage Network)	Configure Brocade 6510 Fabric Network (Block based Storage Network)
Configure storage array infrastructure network	Configure storage array and vSphere host infrastructure networking as specified in the solution document.	Prepare and configure the storage array and Install and configure vSphere hosts
Configure the VLANs	Configure private and public VLANs as required.	Configure VLANs

Prepare Brocade Network Infrastructure

The Brocade network switches deployed with the VSPEX solution provide the redundant links for each ESXi host, the storage array, the switch interconnect ports, and the switch uplink ports. This Brocade storage network configuration provides both scalable bandwidth performance and redundancy. The Brocade network solution can be deployed alongside other components of a newly deployed VSPEX solution or as an upgrade for 1 to 10 GbE transition of existing compute and storage VSPEX solution. This network solution has validated levels of performance and high-availability, this section illustrates the network switching capacity listed in Table 6.

Figure 34 and Figure 35 show a sample redundant Brocade storage network infrastructure for this solution. The diagrams illustrate the use of redundant switches and links to ensure that there are no single points of failure.

Configure Brocade VDX Ethernet Fabric Network (File based Storage Network)

Figure 34 shows a sample redundant Brocade VDX Ethernet Fabric switch for 10 GbE network between compute and storage. The diagram illustrates the use of redundant switches with 10GbE/40GbE links to ensure that no single points of failure exist in the NFS based storage network connectivity.

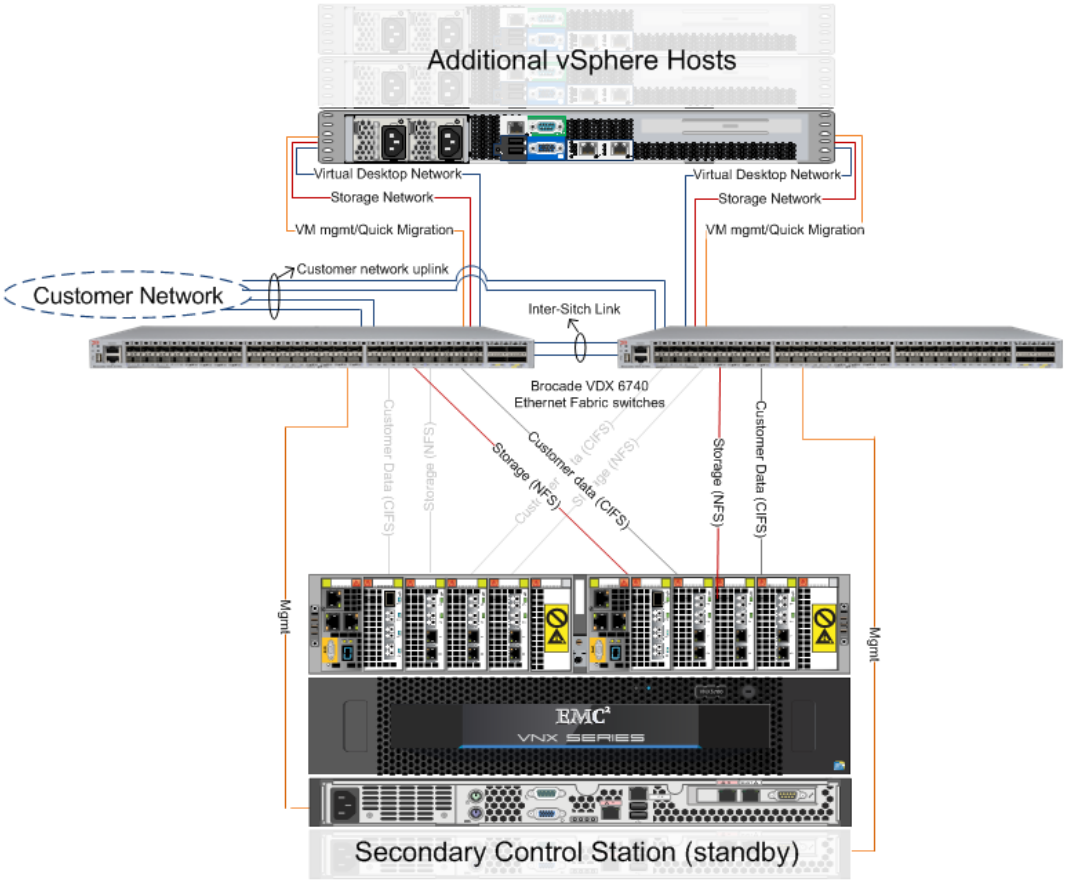

Figure 34. Sample Ethernet network architecture

Note: Ensure there are adequate switch ports between the file based attached storage array & ESXi hosts, and ports to existing customer infrastructure.

Note: Use a minimum of two VLANs for:
- Storage networking (NFS) and vMotion.
- Virtual machine networking and ESXi management (These are customer- facing networks. Separate them if required.)

Note: The Brocade VDX Ethernet Fabric switch provide supports converged network for customers needing FCoE or iSCSI block based storage network attached storage as well.

- Use existing infrastructure that meets the requirements for customer infrastructure and management networks. In this deployment, we are using VLAN 30 for Live Migration traffic, VLAN 20 for Storage traffic and VLAN 10 for Management.

Configure Brocade 6510 Fabric Network (Block based Storage Network)

The infrastructure Figure 35 shows a sample redundant Brocade 6510 Fibre Channel Fabric (FC) switch infrastructure for block based storage network between compute and storage array. The diagram illustrates the use of redundant switches and links to ensure that no single point of failure exist in the network connectivity. The Brocade 6510 switch only supports the FC protocol. The Brocade 6510 FC switches are validated for the FC protocol option.

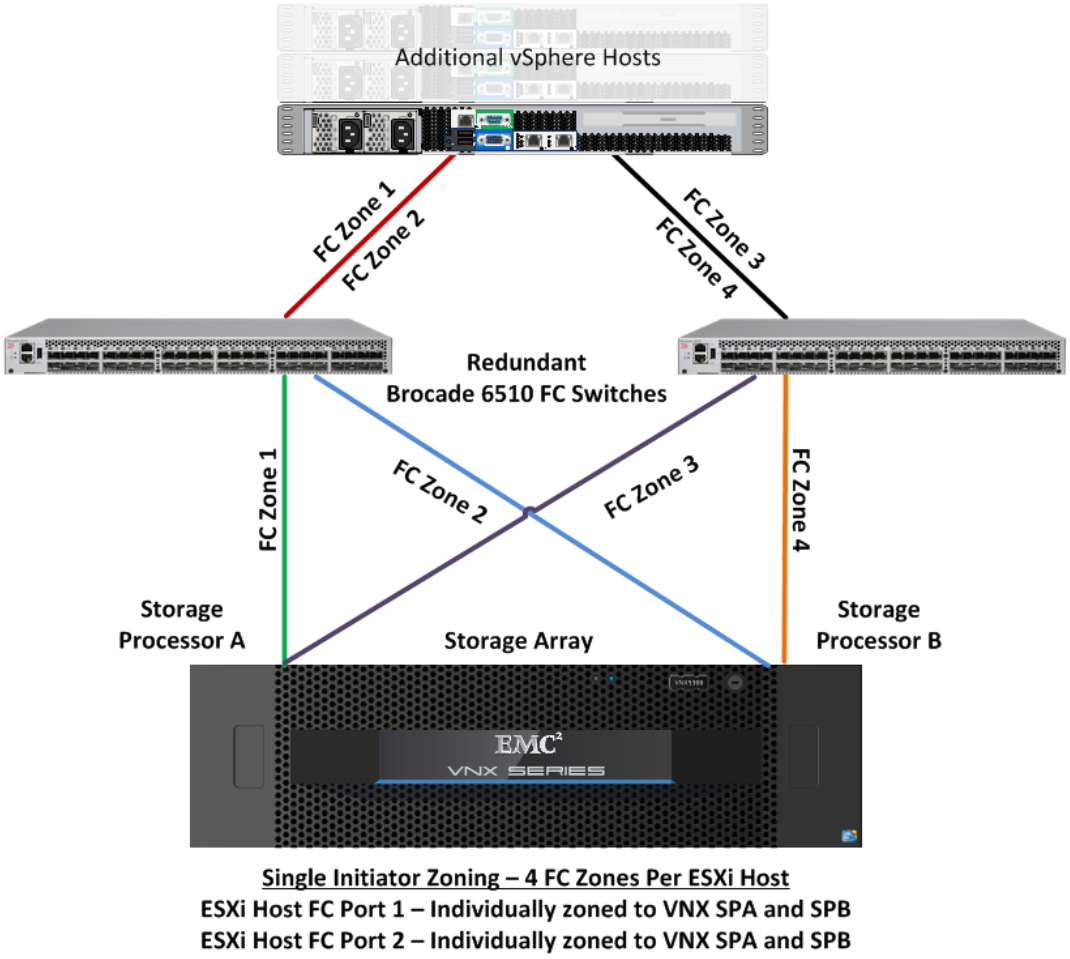

Figure 35. Sample network architecture – Block storage

Configure VLANs

Ensure adequate switch ports for the storage array and vSphere hosts that are configured with a minimum of three VLANs for:

- Management traffic
- NFS networking (private network).
- VMware vMotion (vMotion) (private network).

Complete network cabling

Ensure the following:

- Connect Brocade switch ports to all servers, storage arrays, inter-switch links (ISLs)and uplinks.
- All servers and switch uplinks plug into separate switching infrastructures and have redundant connections.
- Complete connections to the existing customer network.

Note: Brocade switches have Installation Guides providing instructions on racking, cabling, and powering that you can refer to for details.

Note: At this point, the new equipment is being connected to the existing customer network. Ensure that unforeseen interactions do not cause service issues when you connect the new equipment to existing customer infrastructure network.

Configure Brocade VDX 6740 Switch (File Storage)

This section describes Brocade VDX switch configuration procedures for file storage provisioning with VMware. The Brocade VDX switches provide for infrastructure connectivity between ESXi servers, existing customer network, and NFS attached VNX storage as described in the following sections for this VSPEX solution. In this deployment, it is assumed that this new equipment is being connected to the existing customer network and potentially existing compute servers with either 1 GbE or 10 GbE attached NICs.

VSPEX with the Brocade VDX 6740 switches for 10GbE attached ESXi servers are enabled with VCS Fabric Technology which has the following salient features:

- It is an Ethernet fabric switched network. The Ethernet fabric utilizes an emerging standard called Transparent Interconnection of Lots of Links (TRILL) as the underlying technology.
- All switches automatically know about each other and all connected physical and logical devices.
- All paths in the fabric are available. Traffic is always distributed across equal-cost paths. Traffic from the source to the destination can travel across multiple equal cost paths.
- Traffic always travels across the shortest path in the fabric.
- If a single link fails, traffic is automatically switched to other available paths. If one of the links in Active Path #1 goes down, traffic is seamlessly **switched across** Active Path #2.
- Spanning Tree Protocol (STP) is not necessary because the Ethernet fabric itself is loop free and to connected servers, devices, and the rest of the network appears as a single logical switch.
- The fabric is self-forming. When two Brocade VCS Fabric mode-enabled switches are connected, the fabric is automatically created and the switches discover the common fabric configuration.
- The fabric is masterless. No single switch stores configuration information or controls fabric operations. Any switch can fail or be removed without causing disruptive fabric downtime or delayed traffic.

The fabric is aware of all members, devices, and virtual machines (VMs). If the VM moves from one Brocade VCS Fabric port to another Brocade VCS Fabric port in the same fabric, the port-profile is automatically moved to the new port, leveraging Brocade's Automatic Migration of Port Profiles (AMPP) feature.

- All switches in an Ethernet fabric can be managed as if they were a single logical chassis. To the rest of the network, the fabric looks no different from any other Layer 2 switch (Logical Chassis feature).

Brocade VDX switches are available in both port side exhaust and port side intake configurations. Depending upon the hot-aisle, cold-aisle considerations you should choose the appropriate airflow models for your deployment. For more information refer to the Brocade VDX 6740 Hardware Reference Manual as provided in Appendix C.

Listed below is the procedure required to deploy the Brocade VDX 6740 switches with VCS Fabric Technology in VSPEX End-User Computing deployment for up to 2000 Virtual Machines.

Brocade VDX 6740 Configuration Steps

Brocade VDX Configuration Steps
Step 1: Verify and Apply Brocade VDX NOS Licenses
Step 2: Configure Logical Chassis VCS ID and RBridge ID
Step 3: Assign Switch Name
Step 4: Brocade VCS Fabric ISL Port Configuration
Step 5: Create vLAG for ESXi Host
Step 6: vCenter Integration for AMPP
Step 7: Create the vLAG for VNX ports
Step 8: Connecting the VCS Fabric to existing Infrastructure through Uplinks
Step 9 Configure MTU and Jumbo Frames (for NFS)
Step 10: Enable Flow Control Support
Step 11- Auto QOS for NAS
Refer to Appendix C for related documents.

Step 1: Verify and Apply Brocade VDX NOS Licenses

Before starting the switch configurations, make sure you have the required licenses for the VDX 6740 Switches available. With NOS version 4.1.0 or later Brocade VCS Fabric license is built into the code so you will only require port upgrade licenses depending on the number of port density required in the setup. For this deployment, we will be assuming that **48x10GbE** ports are activated on the base Brocade 6740s in the setup and we will be applying one 40GbE Port Upgrade License on each switch which will enable two 40GbE ports on each box. We will use these ports for Inter Switch Links (ISLs) between the two VDXs.

A. Displaying the Switch License ID

The switch license ID identifies the switch for which the license is valid. You will need the switch license ID when you activate a license key, if applicable.

To display the switch license ID, enter the show license id command in the privileged EXEC mode, as shown.

```
Sw0# show license id
Rbridge-Id                          License ID
================================================
  1                     10:00:00:27:F8:BB:7E:85
```

B. Applying the Licenses to the switches

Once you have the **40G Port Upgrade** license strings generated from Brocade's Licensing portal for both the switches, you can apply them to the switches, as shown.

```
sw0# license add licStr "*B
Iflp5mb:NvYn,E4pLcOsVJfqrXDeeu9nMwqM2bQhqtf96TiqVORiWThxA:qsmQ8L3f
IB0tJbTsSuRW,Sfl60zkfbeI2IQiEjHjZFgVb1HLbwLWd3l2JXaDtvcR8DxwiC:wfU
#"

2014/03/07-03:40:27, [SEC-3051], 552,, INFO, sw0, The license key
*B
Iflp5mb:NvYn,E4pLcOsVJfqrXDeeu9nMwqM2bQhqtf96TiqVORiWThxA:qsmQ8L3f
IB0tJbTsSuRW,Sfl60zkfbeI2IQiEjHjZFgVb1HLbwLWd3l2JXaDtvcR8DxwiC:wfU
# is Added.

License Added [*B
Iflp5mb:NvYn,E4pLcOsVJfqrXDeeu9npwqM2bQhqtf96TiqVORiWThxA:qsmQ8L3f
IB0tJbTsSuRW,Sfl60zkfbeI2IQiEjHjZFgVb1HLbwLWd3l2JXaDtvcR8DxwiC:wfU
#]

For license change to take effect, it may be necessary to enable
ports...
```

As noted in the switch output, please note that you may have to enable ports for the licenses to take effect. You can do that by doing **no shut** on the interfaces you are using. The 40GbE ports can also be used in breakout mode as four 10GbE ports. For details, please refer to the *Network OS Administration Guide, v4.1.0* to configure them.

C. Displaying Licenses on the switches

You can display installed licenses with the **show license** command. The following example displays a Brocade VDX 6740 licensed for full port density of 48 ports and two 40Gbe QSFP ports. This configuration does not include FCoE features.

```
sw0# show license

rbridge-id: 1
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
       10G Port Upgrade license
       Feature name:PORT_10G_UPGRADE
       License is valid
       Capacity: 24
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
       40G Port Upgrade license
       Feature name:PORT_40G_UPGRADE
       License is valid
       Capacity: 2
```

Refer to Network OS Administrator's Guide Supporting Network OS v3.0.1 in Appendix C for additional licensing related information.

Step 2: Configure Logical Chassis VCS ID and RBridge ID

When VCS is deployed as a Logical Chassis, it can be managed by a single Virtual IP and configuration changes are automatically saved across all switches in the fabric. RBridge ID is a unique identifier for an RBridge (physical switch in a VCS fabric) and VCS ID is a unique identifier for a VCS fabric. The factory default VCS ID is 1. All switches in a VCS fabric must have the same VCS ID. The default is set to 1 on each VDX switch so it doesn't need to be changed in a one-cluster implementation. The RBridge ID is also set to 1 by default on each VDX switch, but if more than one switch is to be added to the fabric then each switch needs its own unique ID as in this implementation.

In this deployment VCS ID 1 is assigned on all VDXs and RBridge IDs are assigned as per the Deployment Topology in Figure 39. In the following example we will show configuration for Logical Chassis with RB21 as primary. Other Rbridges can be configured in a similar manner.

The value range for RBridge ID is 1-239.
The value range for VCS ID is 1-8192.

In Privileged EXEC mode, enter the ***vcs*** command with options to set the VCS ID, RBridge ID and enable logical chassis mode for the switch. After you execute the below command you are asked if you want to apply the default configuration and reboot the switch; answer 'Yes'.

```
sw0# vcs vcsid 1 rbridge-id 21 logical-chassis enable
This operation will perform a VCS cluster mode transition for this
local node with new parameter settings. This will change the
configuration to default and reboot the switch. Do you want to
continue? [y/n]:
Y
```

Note: To create a Logical Chassis cluster, the user needs to perform the above steps on every VDX in the VCS fabric, changing only the RBridge ID each time, based on Figure 45. Any global and local configuration changes now made are distributed automatically to all nodes in the logical chassis cluster.

You can enter the configuration mode for any VDX in the cluster from the cluster principal node using their respective Rbridge ID.

Optionally, the cluster can also be managed by a Virtual IP (in Logical Chassis and Fabric Cluster mode only) which is tied to the principal node/switch in the cluster. The management interface of the principal switch can be accessed by means of this Virtual IP address, as shown:

```
sw0(config)# vcs virtual ip address 10.246.54.150
```

In the above example, the entire fabric can be managed with one Virtual IP- 10.246.54.150.

Note: For details on Logical Chassis and Virtual Cluster IP, please refer to the *Network OS Administration Guide, v4.1.0*.

Step 3: Assign Switch Name

Every switch is assigned the default host name of "sw0," but must be changed for easy recognition and management using the **switch-attributes** command. Use the **switch-attributes** command to set host name, as shown:

```
sw0# configure terminal
sw0(config)# switch-attributes 21 host-name BRCD6740-RB21
```

After you have enabled the Logical Chassis mode and assigned switch names on each node in the cluster, run the ***show vcs*** command to determine which node has been assigned as the cluster principal node. This node can be used to configure the entire VCS fabric. The arrow (>) denotes the cluster principal node. The asterisk (*) denotes the current logged-in node.

```
BRCD6740-21# show vcs
Config Mode    : Distributed
VCS Mode       : Logical Chassis
VCS ID         : 1
VCS GUID       : 34f262b4-e64f-4a18-a986-a767d389803e
Total Number of Nodes           : 2
Rbridge-Id  WWN  Management IP VCS Status Fabric Status HostName
-----------------------------------------------------------
21 >10:00:00:27:F8:BB:94:18* 10.254.5.44 Online Online BRCD6740-21
22  10:00:00:27:F8:BB:7E:85  10.254.5.43 Online Online BRCD6740-22
….
….
<truncated output>
```

Step 4: Brocade VCS Fabric ISL Port Configuration

The VDX platform comes preconfigured with a default port configuration that enables ISL and Trunking for easy and automatic VCS fabric formation. However, for edge port devices the port configuration requires editing to accommodate specific connections.

The interface format is: rbridge id/slot/port number

For example: 21/0/49

The default port configuration for the 40Gb ports can be seen with the **show running-configuration command**, as shown:

```
BRCD6740-21# show running-config interface FortyGigabitEthernet
21/0/49
interface FortyGigabitEthernet 21/0/49
 fabric isl enable
 fabric trunk enable
 no shutdown
!
<truncated output>
```

There are two types of ports in a VCS fabric, ISL ports, and the edge ports. The ISL port connects VCS fabric switches whereas edge ports connect to end devices or non-VCS Fabric mode switches or routers.

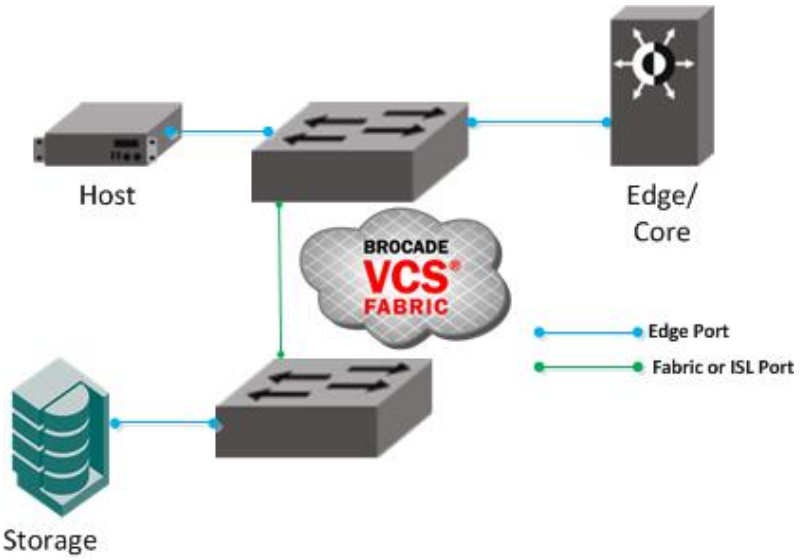

Figure 36. Port types

Fabric ISLs and Trunks

Brocade ISLs connect VDX switches in VCS mode. All ISL ports connected to the same neighbor VDX switch attempt to form a trunk. Trunk formation requires that all ports between the switches are set to the same speed and are part of the same port group. For redundancy purposes the recommendation is to have at least two trunks between two VDXs, but the actual number of trunks required may vary depending on customer I/O/Bandwidth and subscription ratio requirements. The maximum number of ports allowed per trunk group is sixteen on Brocade VDX 6740 and eight in Brocade VDX 6720 and 6730.

In a deployment with the Brocade VDX 6740s, the 10GbE or 40GbE ports can be used for ISLs. In this example configuration, we are using the 40 GbE ports to form two ISLs each which guarantee frame-based load balancing across the ISLs.

Shown below are the port groups for the VDX 6740 and 6740T platforms.

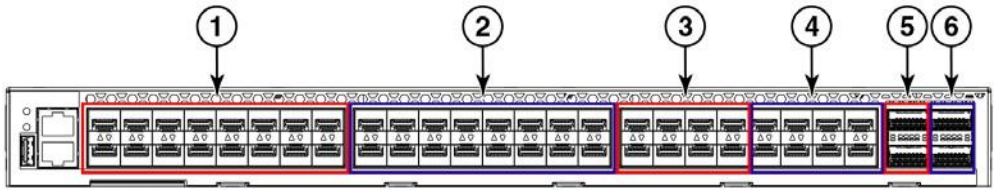

1	Trunk Group 1 - 1/10 GbE SFP ports 1-16	4	Trunk group 4 - 1/10 GbE SFP ports 41-48
2	Trunk Group 2 - 1/10 GbE SFP ports 17-32	5	Trunk Group 3A - 40 GbE QSFP ports 49-50
3	Trunk Group 3 - 1/10 GbE SFP ports 33-40	6	Trunk Group 4A - 40 GbE QSFP ports 51-52

Figure 37. Port Groups of the VDX 6740

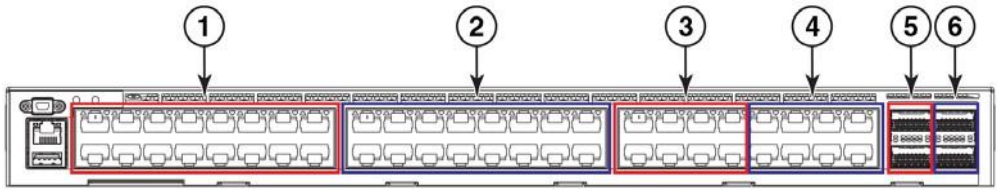

1	Trunk Group 1 - 1/10 GbE BaseT ports 1-16	4	Trunk Group 4 - 1/10 GbE BaseT ports 41-48
2	Trunk Group 2 - 1/10 GbE BaseT ports 17-32	5	Trunk Group 3A - 40 GbE QSFP ports 49-50
3	Trunk Group 3 - 1/10 GbE BaseT ports 33-40	6	Trunk Group 4A - 40 GbE QSFP ports 51-52

Figure 38. Port Groups of the VDX 6740T and Brocade VDX 6740T-1G

Note: On the Brocade VDX 6740, ports in groups 3 and 3A, as well as port groups 4 and 4A, can be trunked together only when the 40 GbE QSFP ports are configured in breakout mode. On the Brocade VDX 6740T/6740T-1G model, this trunking is not allowed.

For more information about Brocade trunking, refer to the *Brocade Network OS Administrator's Guide*, v4.1.0

You can use the ***fabric isl enable***, ***fabric trunk enable***, ***no fabric isl enable***, and ***no fabric trunk enable*** to toggle the ports which are part of a trunked ISL, if needed. The following example shows the running configuration of an ISL port on RB21-

```
BRCD6740-21# show running-config interface FortyGigabitEthernet
21/0/49
interface FortyGigabitEthernet 21/0/49
 fabric isl enable
 fabric trunk enable
 no shutdown
!
```

You can also verify ISL configurations using the **show fabric isl** or **show fabric trunk** commands on RB21, as shown:

```
BRCD6740-21# show fabric isl
Rbridge-id: 21   #ISLs: 2

Src    Src       Nbr    Nbr        Nbr-WWN    BW    Trunk  Nbr-Name
Index  Interface Index  Interface
------------------------------------------------------------------------
0   Fo   21/0/49   0  Fo  22/0/49  10:00:00:27:F8:BB:7E:85 40G Yes
"BRCD6740-22"
2   Fo 21/0/51 2  Fo 22/0/51  10:00:00:27:F8:BB:7E:85   40G   Yes
"BRCD6740-22"
```

```
BRCD6740-21# show fabric trunk
Rbridge-id: 21
Trunk    Src      Source    Nbr    Nbr
Group   Index   Interface  Index   Interface      Nbr-WWN
------------------------------------------------------------------------
 1   0    Fo      21/0/49  0   Fo 22/0/49  10:00:00:27:F8:BB:7E:85
 2   2    Fo      21/0/51  2   Fo 22/0/51  10:00:00:27:F8:BB:7E:85
```

Step 5: Create vLAG for ESXi Hosts

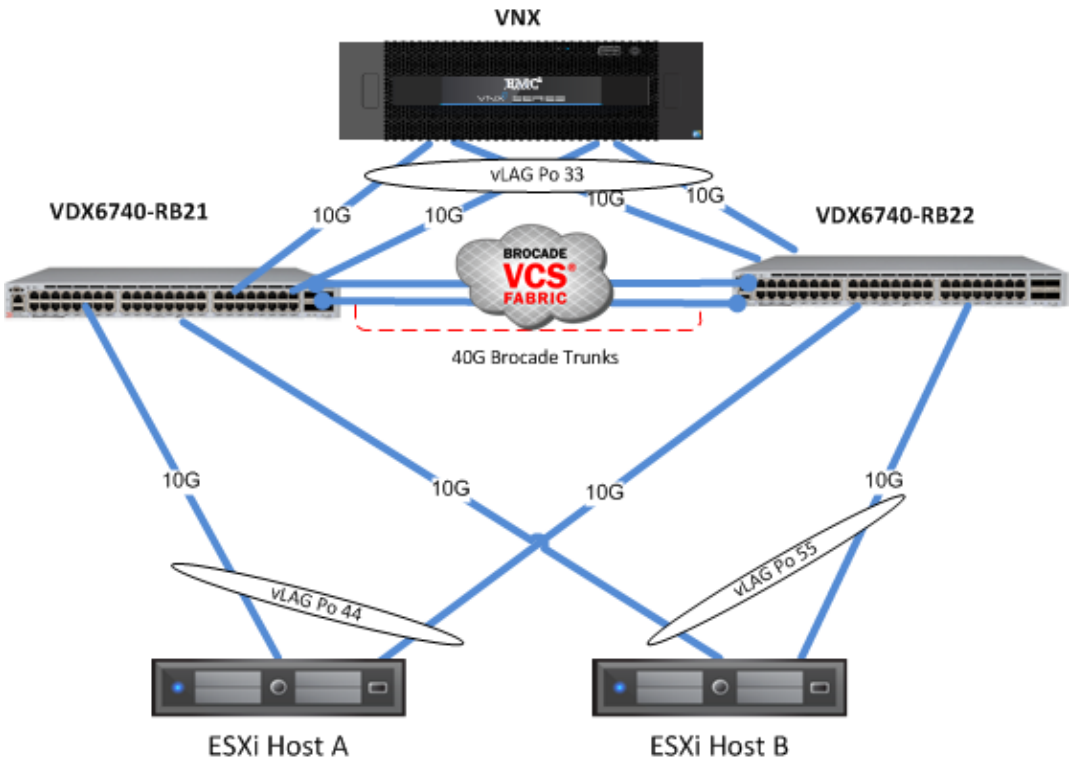

Figure 39. VDX 6740 vLAG for ESXi hosts

Create a port-channel

When creating a port channel interface on both Brocade VDX 6740 switches- RB21 and RB22 note that the port channel number should be the same on both the VDXs, as shown below. Also note that because this solution utilizes vCenter integration, the "switchport" command will not be used and the port will not be configure as a trunk port. This is handled by the vCenter integration.

Configuring Port-channel 44 between Host and VDX switches

Configuration on RB21

```
BRCD6740-RB21(config)# interface port-channel 44
BRCD6740-RB21(config-Port-channel-44)# mtu 9216
BRCD6740-RB21(config-Port-channel-44)# no shutdown
BRCD6740-RB21(config-Port-channel-44)# interface
TenGigabitEthernet 21/0/21
BRCD6740-RB21(conf-if-gi-21/0/21)# channel-group 44 mode on
```

Note: The mode "on" configures the interface as a static vLAG.

Configuration on RB22

```
BRCD6740-RB22# configure terminal
BRCD6740-RB22(config)# interface port-channel 44
BRCD6740-RB22(config-Port-channel-44)# mtu 9216
BRCD6740-RB22(config-Port-channel-44)# no shutdown
BRCD6740-RB22(config-Port-channel-44)# interface
TenGigabitEthernet 22/0/21
BRCD6740-RB22(conf-if-te-22/0/21)# channel-group 44 mode on
```

Configuring Port-channel 55 between Host and VDX switches

Configuration on RB21

```
BRCD6740-RB21# configure terminal
BRCD6740-RB21(config)# interface port-channel 55
BRCD6740-RB21(config-Port-channel-55)# mtu 9216
BRCD6740-RB21(config-Port-channel-55)# no shutdown
BRCD6740-RB21(config-Port-channel-55)# interface
TenGigabitEthernet 21/0/22
BRCD6740-RB21(conf-if-te-21/0/22)# channel-group 55 mode on
```

Configuration on RB22

```
BRCD6740-RB22# configure terminal
BRCD6740-RB22(config)# interface port-channel 55
BRCD6740-RB22(config-Port-channel-55)# mtu 9216
BRCD6740-RB22(config-Port-channel-55)# no shutdown
BRCD6740-RB22(config-Port-channel-55)# interface
TenGigabitEthernet 22/0/22
BRCD6740-RB22(conf-if-te-22/0/22)# channel-group 55 mode on
```

Repeat steps for configuring Port-channel 55 between ESXi Host Band VDX switches, based on Figure 45.

Step 6: vCenter Integration for AMPP

Brocade AMPP (Automatic Migration of Port Profiles) technology enhances network-side virtual machine migration by allowing VM migration across physical switches, switch ports, and collision domains. In traditional networks, port-migration tasks usually require manual configuration changes as VM migration across physical server and switches can result in non-symmetrical network policies. Port setting information must be identical at the destination switch and port.

Brocade VCS Fabrics support automatically moving the port profile in synchronization with a VM moving to a different physical server. This allows VMs to be migrated without the need for network ports to be manually configured on the destination switch.

Port Profile

A port profile contains the entire configuration needed for a VM to gain access to the LAN. The contents of a port profile can be LAN configuration, FCoE configuration, or both.

Specifically, the port profile will contain the VLAN rules, QoS rules and Security ACLs. Depending on the hypervisor there are two ways to configure port profiles - manually or automatically. VDX switches support VMware vCenter integration and this is the preferred method.

vCenter Integration

Note: Before vCenter Integration please make sure required VLAN configuration has been completed on ESXi Hosts.

Brocade best practice is to separate network traffic on different VLANs as shown:

VLAN Name	VLAN ID	VLAN Description
Storage VLAN	20	This VLAN is for NFS traffic
Cluster VLAN	30	This VLAN is for cluster live migration
Management VLAN	10	Management VLAN

We are using VLAN 20 for storage (NFS) traffic in this deployment and hence that needs to be configured in the Port Group properties.

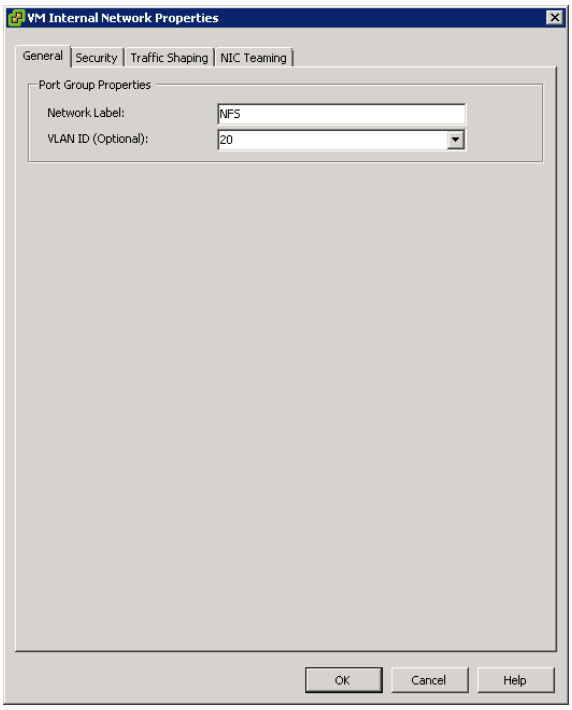

Figure 40. VM Internal Network Properties

VDX 6740 switches support VMware vCenter integration, which provides AMPP automation. NOS v4.1.0 supports vCenter 5.0.

- Automatically creates AMPP port-profiles from VM port groups.
- Automatically creates VLANs.
- Automatically creates association of VMs to port groups.
- Automatically configures port-profile modes on ports.

VDX switches discover and monitor the following vCenter inventory:

- ESX hosts
- Physical network adapters (pNICs)
- Virtual network adapters (vNICs)
- Virtual standard switches (vSwitch)
- Virtual machines (VMs)
- Distributed virtual switch (dvSwitch)
- Distributed virtual port groups

vCenter Integration Process Overview

1. Configure CDP on ESXi hosts.
2. Configure Brocade VDX switch with vCenter access information and credentials.
3. Brocade VDX switch discovers virtual infrastructure assets.

4. VCS fabric will automatically configure corresponding objects including:
   - Port-profiles and VLAN creation
   - MAC address association to port-profiles
   - Port, LAGs, and vLAGs automatically put into profile mode based on ESX host connectivity.
5. VCS fabric is ready for VM movements.

vCenter Integration Configuration Steps

Enable CDP

In order for an Ethernet Fabric to detect the ESX/ESXi hosts, Cisco Discovery Protocol (CDP) must be enabled on all virtual switches (vSwitches) and distributed vSwitches (dvSwitches) in the vCenter Inventory. Each VDX switch in the fabric listens for CDP packets from the ESX hosts on the switch ports. For more information, refer to the VMware KB article 1003885.

Enabling CDP on vSwitches

Login as root to the ESX/ESXi Host. Use the following command to verify the current CDP settings

```
[root@server root]# esxcfg-vswitch -b vSwitch1
```

Use the following command to enable CDP for a given virtual switch. Possible values here are advertise or both.

```
[root@server root]# esxcfg-vswitch -B both vSwitch1
```

Enabling CDP on dvSwitches

1. Connect to the vCenter server using the vSphere Client.
2. In the vCenter Server home page, click **Networking**.
3. Right-click the distributed virtual switches (dvSwitches) and click **Edit Settings**.
4. Select Advanced under **Properties**.
5. Use the check box and the drop-down list to change the **CDP settings**.

Adding and activating vCenter

```
BRCD6740(config)# vcenter production url https://10.20.40.130
username administrator password pass
```

Note: In this example "production" is the name chosen for the vcenter server name.

```
BRCD6740 (config)# vcenter production activate
```

Note: By default, the vCenter server only accepts https connection requests.

Verify vCenter Integration Status

```
BRCD6740-RB21# show vnetwork vcenter status
vCenter      Start                            Elapsed (sec)    Status
-----------------------------------------------------------------------
production 2014-03-09 06:12:43 17       In progress
```

In progress indicates discovery is taking place.
Success will show when it is complete.

Note: Allow at least 30 seconds for the vCenter discovery to complete and show as "Success."

Discovery timer interval

By default, Network Operating System (NOS) queries the vCenter updates every three minutes. NOS detects changes and automatically reconfigures the Ethernet Fabric during the next periodic rediscovery attempt. Detectible changes include any modification to virtual assets, such as adding or deleting virtual machines (VMs), or changing VLANs.

Use the vcenter MYVC interval command to manually change the default timer interval value to suit the individual environment needs.

```
BRCD6740(config)# vcenter MYVC interval ?
Possible completions:
<NUMBER:0-1440> Timer Interval in Minutes (default = 3)
```

Note: Best practice is to keep the discovery timer interval value at default.

User-triggered vCenter discovery

Use the vnetwork vcenter command to manually trigger a vCenter discovery.

```
BRCD6740# vnetwork vcenter MYVC discover
```

Commands to verify AMPP and vCenter integration

Refer to the Network OS Command Reference for detailed information about the show vnetwork command.

Commands that may be useful when monitoring AMPP and vCenter integration include the following.

dvpgs - Displays discovered distributed virtual port groups.

Dvs - Displays discovered distributed virtual switches.

Hosts - Displays discovered hosts.

Pgs - Displays discovered standard port groups.

vcenter status - Displays configured vCenter status.

Vmpolicy - Displays the following network policies on the Brocade VDX switch: associated media access control (MAC) address, virtual machine, (dv) port group, and the associated port profile.

vms - Displays discovered virtual machines (VMs).

vss - Displays discovered standard virtual switches.

Commands to monitor AMPP

```
BRCD6740-RB21# show mac-address-table port-profile
Legend: Untagged(U), Tagged (T), Not Forwardable(NF) and
Conflict(C)
VlanId  Mac-address    Type     State  Port-Profile Ports
1       005a.8402.0007 Dynamic Active Profiled(T)  Te 21/0/21
1       005b.8402.0001 Dynamic Active Profiled(T)  Te 21/0/21
1       005c.8402.0001 Dynamic Active Profiled(T)  Te 21/0/21
```

```
BRCD6740-RB21# show running-config port-profile
port-profile default
vlan-profile
switchport
switchport mode trunk
switchport trunk allowed vlan all
!
!
port-profile vm_kernel
vlan-profile
switchport
switchport mode access
switchport access vlan 1
```

```
BRCD6740-RB21# show port-profile
port-profile default
ppid 0
vlan-profile
switchport
switchport mode trunk
switchport trunk allowed vlan all
port-profile vm_kernel
ppid 1
vlan-profile
switchport
switchport mode access
switchport access vlan 1
```

```
BRCD6740-RB21# show port-profile status activated
Port-Profile      PPID     Activated Associated MAC Interface
auto-dvPortGroup 1    Yes               None     None
auto-dvPortGroup2 2   Yes              None      None
auto-dvPortGroup3 3   Yes              None      None
auto-dvPortGroup_4_0 4 Yes    0050.567e.98b0    None
auto-dvPortGroup_vlag 5 Yes   0050.5678.eaed    None
auto-for_nfs 6          Yes   0050.5673.85f9    None
```

```
BRCD6740-RB21# show port-profile status associated
Port-Profile PPID Activated Associated MAC Interface
auto-dvPortGroup_4_0 4 Yes 0050.567e.98b0 None
auto-dvPortGroup_vlag 5 Yes 0050.5678.eaed None
auto-for_nfs 6 Yes 0050.5673.85f9 None
```

Step 7: Create the vLAG for VNX ports

Some current storage arrays like EMC's VNX 5400 support LACP-based dynamic LAGs, so in order to provide link and node level redundancy, dynamic LACP based vLAGS can be configured on the Brocade VDX switches.

Note: In some port-channel configurations, depending on the storage ports (1G or 10G), the speed on the port-channel might need to be set manually on the VDX 6740 as shown in the following example:

```
BRCD6740# configure terminal
BRCD6740(config)# interface Port-channel 33
BRCD6740(config-Port-channel-33)# speed
[1000,10000,40000] (1000): 10000
BRCD6740(config-Port-channel-33)#
```

To configure dynamic vLAGS on each Brocade VDX 6740 switch interface, use the following steps-

1. Configure vLAG Port-channel Interface on BRCD6740-RB21 for VNX(enabled for storage VLAN 20)

```
BRCD6740-RB21# configure terminal
BRCD6740-RB21(config)# interface Port-channel 33
BRCD6740-RB21(config-Port-channel-33)# mtu 9216
BRCD6740-RB21(config-Port-channel-33)# description VNX-vLAG-33
BRCD6740-RB21(config-Port-channel-33)# switchport
BRCD6740-RB21(config-Port-channel-33)# switchport mode trunk
BRCD6740-RB21(config-Port-channel-33)# switchport trunk allowed
vlan 20
```

2. Configure Interface TenGigabitEthernet 21/0/23 and 21/0/24 on BRCD6740-RB21 for Port-Channel 33.

```
BRCD6740-RB21# configure terminal
BRCD6740-RB21(config)# interface TenGigabitEthernet 21/0/23
BRCD6740-RB21(conf-if-te-21/0/23)# description VNX-SPA-fxg-1-0
BRCD6740-RB21(conf-if-te-21/0/23)# channel-group 33 mode active
type standard
BRCD6740-RB21(conf-if-te-21/0/23)# lacp timeout long
BRCD6740-RB21(conf-if-te-21/0/23)# no shutdown
```

```
BRCD6740-RB21# configure terminal
BRCD6740-RB21(config)# interface TenGigabitEthernet 21/0/24
BRCD6740-RB21(conf-if-te-21/0/24)# description VNX-SPA-fxg-1-1
BRCD6740-RB21(conf-if-te-21/0/24)# channel-group 33 mode active
type standard
BRCD6740-RB21(conf-if-te-21/0/24)# lacp timeout long
BRCD6740-RB21(conf-if-te-21/0/24)# no shutdown
```

3. Repeat above steps 1-2 on Logical Chassis' Principal Node Rbridge 21 to configure Port-Channel 33 on Rbridge 22's interfaces 22/0/23 and 22/0/24((enabled for storage VLAN 20)going to SPB on the VNX.

4. Validate vLAG Port-channel Interface on BRCD6740-RB21 and BRCD6740-RB22 to VNX.

```
BRCD6740-RB21# show interface Port-channel 33
Port-channel 33 is up, line protocol is up
Hardware is AGGREGATE, address is 0005.338c.adee
Current address is 0005.338c.adee
Description: VNX-vLAG-33
Interface index (ifindex) is 672088673
Minimum number of links to bring Port-channel up is 1
MTU 9216 bytes
LineSpeed Actual : 10000 Mbit
Allowed Member Speed : 10000 Mbit
```

Step 8: Connecting the VCS Fabric to existing Infrastructure through Uplinks

Brocade VDX 6740 switches can be uplinked to be accessible from customer's existing network infrastructure. On VDX 6740 platforms, the user will need to use 40GbE or 10GbE uplinks. The uplink should be configured to match whether or not the customer's network is using tagged or untagged traffic.

The following example can be leveraged as a guideline to connect VCS fabric to existing infrastructure network:

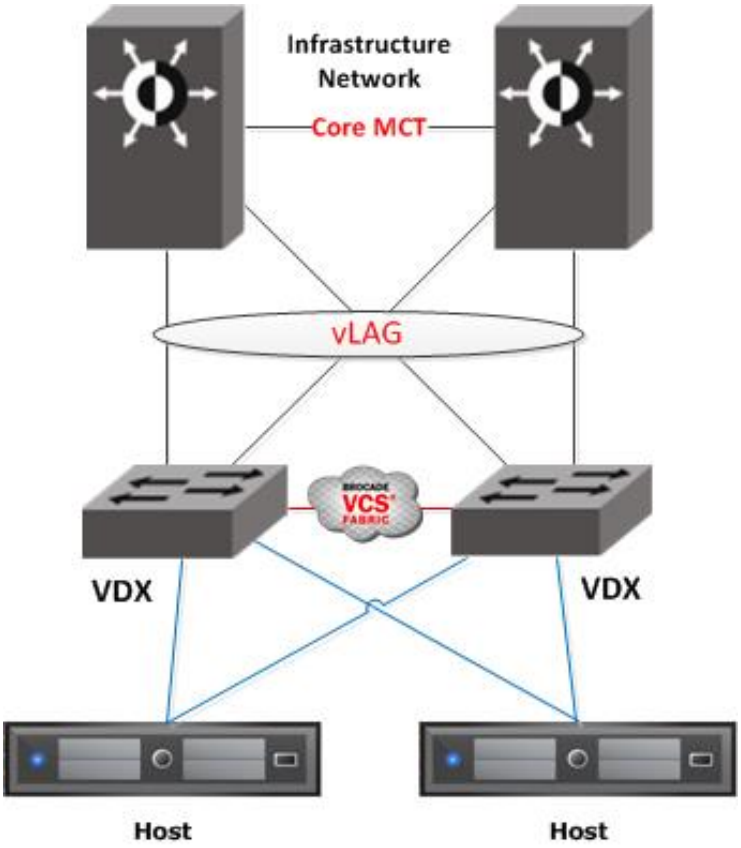

Figure 41. Example VCS/VDX network topology with Infrastructure connectivity

Creating virtual link aggregation groups (vLAGs) to the Infrastructure Network

Create vLAGs from each RBridge to Infrastructure Switches that in turn provide access to resources at the core network.

This example illustrates the configuration for Port-Channel 4 on RB21 and RB22.

1. Use the **switchport** command to configure Port-Channel 4 interface. The following example assigns it to trunk mode and allows all VLANs on the port channel.

```
BRCD6740-RB21(config)# interface port-channel 4
BRCD6740-RB21(config-Port-channel-4)# switchport
BRCD6740-RB21(config-Port-channel-4)# switchport mode trunk
BRCD6740-RB21(config-Port-channel-4)# switchport trunk allowed
vlan all
BRCD6740-RB21(config-Port-channel-4)# no shutdown
```

2. Use the **channel-group** command to configure interfaces as members of a Port-Channel 4 to the infrastructure switches that interface to the core.

```
BRCD6740-RB21(config)# in te 21/0/5
BRCD6740-RB21(conf-if-te-21/0/5)# channel-group 4 mode active type
standard
BRCD6740-RB21(conf-if-te-21/0/5)# in te 21/0/6
BRCD6740-RB21(conf-if-te-21/0/6)# channel-group 4 mode active type
standard
```

3. Repeat above steps 1-2 on Logical Chassis' Principal Node Rbridge 21 to configure Port-Channel 4 on Rbridge 22's interfaces 22/0/5 and 22/0/6.

4. Use the **do show port-chan** command to confirm that the vLAG comes up and is configured correctly.

Note: The LAG must be configured on the MLX MCT as well before the vLAG can become operational.

```
BRCD6740-RB21(config-Port-channel-4)# do show port-chan 4
LACP Aggregator: Po 4 (vLAG)
Aggregator type: Standard
Ignore-split is enabled
Member rbridges:
rbridge-id: 21 (2)
rbridge-id: 22 (2)
Admin Key: 0004 - Oper Key 0004
Partner System ID - 0x0001,01-80-c2-00-00-01
Partner Oper Key 30002
Member ports on rbridge-id 21:
Link: Te 21/0/5 (0x151810000F) sync: 1 *
Link: Te 21/0/6 (0x1518110010) sync: 1
```

```
BRCD6740-RB22(config-Port-channel-4)# do show port-channel 4
LACP Aggregator: Po 4 (vLAG)
Aggregator type: Standard
Ignore-split is enabled
Member rbridges:
rbridge-id: 21 (2)
rbridge-id: 22 (2)
Admin Key: 0004 - Oper Key 0004
Partner System ID - 0x0001,01-80-c2-00-00-01
Partner Oper Key 30002
Member ports on rbridge-id 22:
Link: Te 22/0/5 (0x161810000F) sync: 1
Link: Te 22/0/6 (0x1618110010) sync: 1
```

Step 9 Configure MTU and Jumbo Frames (for NFS)

Brocade VDX Series switches support the transport of jumbo frames. This solution recommends an MTU setting at 9216 (Jumbo frames) for efficient NAS storage and migration traffic. Jumbo frames are enabled by default on the Brocade ISL trunks. However, to accommodate end-to-end jumbo frame support on the network for the edge systems, this feature can be enabled under the vLAG interface. Please note that for end-to-end flow control, Jumbo frames need to be enabled both on the host servers and the storage with the same MTU size of 9216.

Configuring MTU

Note: This must be performed on all RBbridges where a given interface port-channel is located. In this example, interface port-channel 44 is on RBridge 21 and RBridge 22, so we will apply configurations from both RBridge 21 and RBridge 22.

Example to enable Jumbo Frame Support on applicable VDX interfaces for which Jumbo Frame support is required:

```
BRCD6740-RB21# configure terminal
BRCD6740-RB21(config)# interface Port-channel 44
BRCD6740-RB21(config-Port-channel-44)# mtu
(<NUMBER:1522-9216>) (9216): 9216
```

Step 10: Enable Flow Control Support

Ethernet Flow Control is used to prevent dropped frames by slowing traffic at the source end of a link. When a port on a switch or host is not ready to receive more traffic from the source, perhaps due to congestion, it sends pause frames to the source to pause the traffic flow. When the congestion is cleared, the port stops requesting the source to pause traffic flow, and traffic resumes without any frame drop. It is recommended to enable Flow Control on vLAG interfaces towards the VNX on the VDX 6740s, as shown:

Enable QOS Flow Control for both tx and rx on RB21 and RB22

```
BRCD6740-RB21# conf t
BRCD6740-RB21 (config)# interface Port-channel 33
BRCD6740-RB21 (config-Port-channel-33)# qos flowcontrol tx on rx
on
```

Step 11- Auto QOS for NAS

The Auto QoS feature introduced in NOS v4.1.0 automatically classifies traffic based on either a source or a destination IPv4 address. Once the traffic is identified, it is assigned to a separate priority queue. This allows a minimum bandwidth guarantee to be provided to the queue so that the identified traffic is less affected by network traffic congestion than other traffic.

Note: As this command was created primarily to benefit Network Attached Storage devices (NAS). The commands used with this feature use the term tures created primarily too strict requirement that these nodes be actual NAS devices, as Auto QoS will prioritize the traffic for any set of specified IP addresses.

There are four steps to enabling and configuring Auto QoS for NAS:

1. Enable Auto QoS.
2. Set the Auto QoS CoS value.
3. Set the Auto QoS DSCP value.
4. Specify the NAS server IP addresses.

For detailed instructions for setting this feature up, please refer to the *Network OS Administration Guide, v4.1.0*.

Configure Brocade 6510 Switch storage network (Block Storage)

Listed below is the procedure required to deploy the Brocade 6510 Fibre Channel (FC) switches in the EMC® VSPEX™ with Brocade Networking Solutions for END-USER COMPUTING VMware Horizon View 5.3 and VMware vSphere for up to 2,000 Virtual Desktops with block storage network. The Brocade 6510 FC switches provide for infrastructure connectivity between servers and attached VNX storage of the VSPEX solution. At the point of deployment, compute nodes connected FC storage network with either 4, 8, or 16G FC attached HBAs.

- Provide flexibility, simplicity, and enterprise-class functionality in a 48-port switch for virtualized data centers and private cloud architectures
- Enables fast, easy, and cost-effective scaling from 24 to 48 ports using Ports on Demand (PoD) capabilities
- Simplifies deployment with the Brocade EZSwitchSetup wizard
- Accelerates deployment and troubleshooting time with Dynamic Fabric Provisioning (DFP), critical monitoring, and advanced diagnostic features
- Maximizes availability with redundant, hot-pluggable components and non-disruptive software upgrades
- Simplifies server connectivity and SAN scalability by offering dual functionality as either a full-fabric SAN switch or an NPIV-enabled Brocade Access Gateway

In addition, it is important to consider the airflow direction of the switches. Brocade 6510 FC switches are available in both port side exhaust and port side intake configurations. Depending upon the hot-aisle, cold-aisle considerations choose the appropriate airflow. For more information refer to the *Brocade 6510 Hardware Reference Manual* as provided in Appendix C.

All Brocade Fibre Channel Switches have factory defaults listed in Table 27.

Table 27. Brocade switch default settings

Setting	Factory default
Factory Default MGMT IP:	10.77.77.77
Factory Default Subnet:	255.0.0.0
Factory Default Gateway:	0.0.0.0
Factory Default admin/User Password:	password
Factory Default Domain ID:	1
Brocade Switch Management	
	CLI
	Web Tools
	Connectrix Manager

Listed below is the procedure required to deploy the Brocade 6510 FC switches in the VSPEX Private Cloud Solution for up to 500 Virtual Machines.

Table 28. Brocade 6510 FC switch Configuration Steps

Step
Step 1: Initial Switch Configuration
Step 2: Fibre Channel Switch Licensing
Step 3: Zoning Configuration
Step 4: Switch Management and Monitoring
Please see Appendix B for related documents

Step 1: Initial Switch Configuration

Configure Hyper Terminal

1. Connect the serial cable to the serial port on the switch and to an RS-232 serial port on the workstation.
2. Open a terminal emulator application (such as HyperTerminal on a PC) and configure the application as follows

Table 29. Brocade switch default settings

Parameter	Value
Bits per second	9600
Databits	8
Parity	None
Stop bits	1
Flow control	None

Configure IP Address for Management Interface

Switch IP address

You can configure the Brocade 6510 with a static IP address, or you can use a DHCP (Dynamic Host Configuration Protocol) server to set the IP address of the switch. DHCP is enabled by default. The Brocade 6510 supports both IPv4 and IPv6.

Using DHCP to set the IP address

When using DHCP, the Brocade 6510 obtains its IP address, subnet mask, and default gateway address from the DHCP server. The DHCP client can only connect to a DHCP server that is on the same subnet as the switch. If your DHCP server is not on the same subnet as the Brocade 6510, use a static IP address.

Setting a static IP address

1. Log into the switch using the default password, which is password.
2. Use the ipaddrset command to set the Ethernet IP address.

If you are going to use an IPv4 IP address, enter the IP address in dotted decimal notation as prompted. As you enter a value and press Enter for a line in the following example, the next line appears.

For instance, the Ethernet IP Address appears first. When you enter a new IP address and press Enter or simply press Enter accept the existing value, the Ethernet Subnetmask line appears.

In addition to the Ethernet IP address itself, you can set the Ethernet subnet mask, theGateway IP address, and whether to obtain the IP address via Dynamic Host Control Protocol (DHCP) or not.

```
SW6510:admin> ipaddrset
Ethernet IP Address [10.77.77.77]:10.18.226.172
Ethernet Subnetmask [255.255.255.0]:255.255.255.0
Gateway IP Address [0.0.0.0]:10.18.226.1
DHCP [Off]: off
```

If you are going to use an IPv6 address, enter the network information in semicolon-separated notation as a standalone command.

```
SW6510:admin> ipaddrset -ipv6 --add 1080::8:800:200C:417A/64
IP address is being changed...
```

Configure Domain ID and Fabric Parameters

```
RCD-FC-6510:FID128:admin> switchdisable
BRCD-FC-6510:FID128:admin> configure
Configure...
  Fabric parameters (yes, y, no, n): [no] y
    Domain: (1..239) [1] 10
    WWN Based persistent PID (yes, y, no, n): [no]
    Allow XISL Use (yes, y, no, n): [no]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    WAN_TOV: (0..30000) [0]
    MAX_HOPS: (7..19) [7]
    Data field size: (256..2112) [2112]
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0]
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
    BB credit: (1..27) [16]
    Disable FID Check (yes, y, no, n): [no]
  Insistent Domain ID Mode (yes, y, no, n): [no] yes
  Disable Default PortName (yes, y, no, n): [no]
    Edge Hold Time(0 = Low(80ms),1 = Medium(220ms),2 =
High(500ms): [220ms]): (0..2) [1]
  Virtual Channel parameters (yes, y, no, n): [no]
  F-Port login parameters (yes, y, no, n): [no]
  Zoning Operation parameters (yes, y, no, n): [no]
  RSCN Transmission Mode (yes, y, no, n): [no]
  Arbitrated Loop parameters (yes, y, no, n): [no]
  System services (yes, y, no, n): [no]
  Portlog events enable (yes, y, no, n): [no]
  ssl attributes (yes, y, no, n): [no]
  rpcd attributes (yes, y, no, n): [no]
  webtools attributes (yes, y, no, n): [no]
```

Note: The domain ID will be changed. The port level zoning may be affected.

Note: The Domain ID will be changed.

Since Insistent Domain ID Mode is enabled, please ensure that switches in fabric do not have duplicate domain IDs configured, otherwise this may cause switch to segment, if Insistent domain ID is not obtained when fabric re-configures.

```
BRCD-FC-6510:FID128:admin> switchenable
```

Set Switch Name

```
SW6510:FID128:admin> switchname BRCD-FC-6510
Committing configuration...
Done.
```

Verify Domain ID and Switch Name

```
BRCD-FC-6510:FID128:admin> switchshow |more
switchName:     BRCD-FC-6510
switchType:     109.1
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   10
switchId:       fffc0a
switchWwn:      10:00:00:27:f8:61:80:8a
zoning:         OFF
switchBeacon:   OFF
FC Router:      OFF
Allow XISL Use: OFF
LS Attributes:  [FID: 128, Base Switch: No, Default Switch: Yes,
Address Mode 0]
```

Date and Time Setting

The Brocade 6510 maintains the current date and time inside a battery-backed real-time clock (RTC) circuit. Date and time are used for logging events. Switch operation does not depend on the date and time; a Brocade 6510 with an incorrect date and time value still functions properly. However, because the date and time are used for logging, error detection, and troubleshooting, you should set them correctly.

Time Zone, Date and Clock Server can be configured on all Brocade switches.

Time Zone

You can set the time zone for the switch by name. You can also set country, city or time zone parameters.

```
BRCD-FC-6510:FID128:admin> tstimezone --interactive
Please identify a location so that time zone rules can be set
correctly.
Please select a continent or ocean.
 1) Africa
 2) Americas
 3) Antarctica
 4) Arctic Ocean
 5) Asia
 6) Atlantic Ocean
 7) Australia
 8) Europe
 9) Indian Ocean
10) Pacific Ocean
11) none - I want to specify the time zone using the POSIX TZ
format.
Enter number or control-D to quit ?2
Please select a country.
 1) Anguilla                    27) Honduras
 2) Antigua & Barbuda           28) Jamaica
 3) Argentina                   29) Martinique
 4) Aruba                       30) Mexico
 5) Bahamas                     31) Montserrat
 6) Barbados                    32) Netherlands Antilles
 7) Belize                      33) Nicaragua
 8) Bolivia                     34) Panama
 9) Brazil                      35) Paraguay
10) Canada                      36) Peru
11) Cayman Islands              37) Puerto Rico
12) Chile                       38) St Barthelemy
13) Colombia                    39) St Kitts & Nevis
14) Costa Rica                  40) St Lucia
15) Cuba                        41) St Martin (French part)
16) Dominica                    42) St Pierre & Miquelon
17) Dominican Republic          43) St Vincent
18) Ecuador                     44) Suriname
19) El Salvador                 45) Trinidad & Tobago
20) French Guiana               46) Turks & Caicos Is
21) Greenland                   47) United States
22) Grenada                     48) Uruguay
23) Guadeloupe                  49) Venezuela
24) Guatemala                   50) Virgin Islands (UK)
25) Guyana                      51) Virgin Islands (US)
26) Haiti
Enter number or control-D to quit ?47
```

```
Please select one of the following time zone regions.
 1) Eastern Time
 2) Eastern Time - Michigan - most locations
 3) Eastern Time - Kentucky - Louisville area
 4) Eastern Time - Kentucky - Wayne County
 5) Eastern Time - Indiana - most locations
 6) Eastern Time - Indiana - Daviess, Dubois, Knox & Martin
Counties
 7) Eastern Time - Indiana - Starke County
 8) Eastern Time - Indiana - Pulaski County
 9) Eastern Time - Indiana - Crawford County
10) Eastern Time - Indiana - Switzerland County
11) Central Time
12) Central Time - Indiana - Perry County
13) Central Time - Indiana - Pike County
14) Central Time - Michigan - Dickinson, Gogebic, Iron & Menominee
Counties
15) Central Time - North Dakota - Oliver County
16) Central Time - North Dakota - Morton County (except Mandan
area)
17) Mountain Time
18) Mountain Time - south Idaho & east Oregon
19) Mountain Time - Navajo
20) Mountain Standard Time - Arizona
21) Pacific Time
22) Alaska Time
23) Alaska Time - Alaska panhandle
24) Alaska Time - Alaska panhandle neck
25) Alaska Time - west Alaska
26) Aleutian Islands
27) Hawaii
Enter number or control-D to quit ?21
```

```
The following information has been given:

        United States
        Pacific Time

Therefore TZ='America/Los_Angeles' will be used.
Local time is now:      Mon Aug 12 15:04:43 PDT 2013.
Universal Time is now:  Mon Aug 12 22:04:43 UTC 2013.

Is the above information OK?
1) Yes
2) No
Enter number or control-D to quit ?1
System Time Zone change will take effect at next reboot
```

Setting the date

1. Log into the switch using the default password, which is password.
2. Enter the date command, using the following syntax (the double quotation marks are required):

Syntax: **date** "mmddHHMMyy"

The values are:

- mm is the month; valid values are 01 through 12.
- dd is the date; valid values are 01 through 31.
- HH is the hour; valid values are 00 through 23.
- MM is minutes; valid values are 00 through 59.
- yy is the year; valid values are 00 through 99 (values greater than 69 are interpreted as 1970 through 1999, and values less than 70 are interpreted as 2000-2069).

```
switch:admin> date
Fri Sep 29 17:01:48 UTC 2007
switch:admin> date "0927123007"
Thu Sep 27 12:30:00 UTC 2007
switch:admin>
```

Synchronizing local time using NTP

Perform the following steps to synchronize the local time using NTP.

1. Log into the switch using the default password, which is password.
2. Enter the tsClockServer command:

```
switch:admin> tsclockserver "<ntp1;ntp2>"
```

Where ntp1 is the IP address or DNS name of the first NTP server, which the switch must be able to access. The value ntp2 is the name of the second NTP server and is optional. The entire operand "<ntp1;ntp2>" is optional; by default, this value is LOCL, which uses the local clock of the principal or primary switch as the clock server.

```
switch:admin> tsclockserver
LOCL
```

Verify Switch Component Status

```
BRCD-FC-6510:FID128:admin> switchstatusshow
Switch Health Report                          Report time:
08/14/2013 09:19:56 PM
Switch Name:    BRCD-FC-6510
IP address:     10.18.226.172
SwitchState:    HEALTHY
Duration:       218:52
Power supplies monitor  HEALTHY
Temperatures monitor    HEALTHY
Fans monitor            HEALTHY
Flash monitor           HEALTHY
Marginal ports monitor  HEALTHY
Faulty ports monitor    HEALTHY
Missing SFPs monitor    HEALTHY
Error ports monitor     HEALTHY
Fabric Watch is not licensed
Detailed port information is not included
BRCD-FC-6510:FID128:admin>
```

Step 2: FC Switch Licensing

Verify and Install Licenses

Brocade GEN5 Fibre Channel switches come with preinstalled basic licenses required FC operation. The Brocade 6510 provides 48 ports in a single (1U) height switch that enables the creation of very dense fabrics in a relatively small space.

The Brocade 6510 offers Ports on Demand (POD) licensing as well. "Base" models of the switch contain 24 ports, and up to two additional 12-port POD licenses can be purchased.

1. licenseshow (Record License Info) if applicable
2. If POD license needs to be installed on the switch you would require Transaction Key (From License Purchase Paper Pack) and Switch WWN (from wwn –sn or switchshow command output)
3. licenseadd "key" can be used to add the license to the switch.

Obtaining New License Keys

To obtain POD license keys please contact licensekeys@EMC.COM

Step 3: FC Zoning Configuration

Zone Objects

A zone object is any device in a zone, such as:

Physical port number or port index on the switch

Node World Wide Name (N-WWN)

Port World Wide Name (P-WWN)

Zone Schemes

You can establish a zone by identifying a zone objects by using one or more of the following zoning schemes

- **Domain,Index -** All members are specified by Domain ID and Port Number or Domain Index number pair or aliases.
- **World Wide Name (WWN) -** All members are specified only by WWW or Aliases of the WWN. They can be Node or Port version of WWN.
- **Mixed Zoning -** A zone containing members specified by a combination of domain, port or domain, index and wwn .

Configurations of Zones

Following are recommendations for zoning:

Do nsshow to list the WWN of the Host and Storage (Initiator and Target). Record Port WWN

Create the Alias for Device alicreate "Alias, "WWN"

Create the Zone zoneCreate “Zone Name”,”WWN/Alias”

Create the Zone Configuration cfgCreate “cfgName”,”Zone Name”

Save the Zone Configuration cfgsave

Enable the Zone Configuration cfgenable “cfgName”

```
BRCD-FC-6510:FID128:admin> nsshow
{
 Type Pid    COS     PortName               NodeName        TTL(sec)
 N   0a0500; 3;10:00:00:05:33:64:d6:35;20:00:00:05:33:64:d6:35; na
    FC4s: FCP
    PortSymb: [30] "Brocade-825 | 2.3.0.2 |  |  | "
    Fabric Port Name: 20:05:00:27:f8:61:80:8a
    Permanent Port Name: 10:00:00:05:33:64:d6:35
    Port Index: 5
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No

N  0a0a00; 3;50:06:01:6c:36:60:07:c3;50:06:01:60:b6:60:07:c3; na
    FC4s: FCP
    PortSymb: [27] "CLARiiON::::SPB10::FC::::::"
    NodeSymb: [25] "CLARiiON::::SPB::FC::::::"
    Fabric Port Name: 20:05:00:05:1e:02:93:75
```

```
    Permanent Port Name: 50:06:01:6c:36:60:07:c3
    Port Index: 10
    Share Area: No
    Device Shared in Other AD: No
    Redirect: No
    Partial: No
```

The Local Name Server has 2 entries }

Create Alias

```
SW6510:FID128:admin> alicreate
error: Usage: alicreate "arg1", "arg2"
SW6510:FID128:admin> alicreate
"ESX_Host_HBA1_P0","10:00:00:05:33:64:d6:35"
SW6510:FID128:admin> alicreate
"VNX_SPA_P0","50:06:01:60:b6:60:07:c3"
```

Create Zone

```
SW6510:FID128:admin> zonecreate
error: Usage: zonecreate "arg1", "arg2"
SW6510:FID128:admin> zonecreate
"ESX_Host_A","ESX_Host_HBA1_P0;VNX_SPA_P0"
```

Create cfg and add zone to cfg

```
SW6510:FID128:admin> cfgcreate
error: Usage: cfgcreate "arg1", "arg2"
SW6510:FID128:admin> cfgcreate "vspex", "ESX_Host_A"
Save cfg and enable cfg
SW6510:FID128:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.  Any
changes made on the Effective configuration will not take effect
until it is re-enabled. Until the Effective configuration is re-
enabled, merging new switches into the fabric is not recommended
and may cause unpredictable results with the potential of
mismatched Effective Zoning configurations.

Do you want to save Defined zoning configuration only?  (yes, y,
no, n): [no] y
Updating flash ...
SW6510:FID128:admin> cfgenable "vspex"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Do you want to enable 'vspex' configuration  (yes, y, no, n): [no]
y
zone config "vspex" is in effect
Updating flash ...
```

Verify Zone Configuration

```
SW6510:FID128:admin> cfgshow
Defined configuration:
 cfg:    vspex    ESX_Host_A
 zone:   ESX_Host_A
                  ESX_Host_HBA1_P0; VNX_SPA_P0
 alias: ESX_Host_HBA1_P0
                  10:00:00:05:33:64:d6:35
 alias: VNX_SPA_P0
                  50:06:01:60:b6:60:07:c3
Effective configuration:
 cfg:    vspex
 zone:   ESX_Host_A
                  10:00:00:05:33:64:d6:35
                  50:06:01:60:b6:60:07:c3
SW6510:FID128:admin> cfgactvshow
Effective configuration:
 cfg:    vspex
 zone:   ESX_Host_A
                  10:00:00:05:33:64:d6:35
                  50:06:01:60:b6:60:07:c3
```

Follow the above Zoning steps to configure Fabric-B switch.

Step 4: Switch Management and Monitoring

Following table shows a list of commands that can be used to Manage and Monitor Brocade Fibre Channel switches in a production environment.

```
Switch Management        Switchshow
Switch Monitoring
1. porterrshow
2. Portperfshow
3. Portshow
4. errshow
5. Errdump
6. Sfpshow
7. Fanshow
8. Psshow
9. Sensorshow
10. Firmwareshow
11. Fosconfig --show
12. Memshow
13. Portcfgshow
14. Supportsave – to collect switch logs
```

Prepare and configure the storage array

VNX configuration

This section describes how to configure the VNX storage array. In this solution, the VNX series provides NFS or VMware Virtual Machine File System (VMFS) data storage for VMware hosts. Table 30 shows the tasks for the storage configuration.

Table 30. Tasks for storage configuration

Task	Description	Reference
Set up the initial VNX configuration	Configure the IP address information and other key parameters on the VNX.	VNX5400 Unified Installation Guide VNX5600 Unified Installation Guide VNX File and Unified Worksheet Unisphere System Getting Started Guide Vendor's switch configuration guide
Provision storage for VMFS datastores (FC only)	Create FC LUNs that will be presented to the vSphere servers as VMFS datastores hosting the virtual desktops.	
Provision storage for NFS datastores (NFS only)	Create NFS file systems that will be presented to the vSphere servers as NFS datastores hosting the virtual desktops.	
Provision optional storage for user data	Create CIFS file systems that will be used to store roaming user profiles and home directories.	
Provision optional storage for infrastructure virtual machines	Create optional VMFS/NFS datastores to host the SQL server, domain controller, vCenter server, and/or VMware View Connection Server virtual machines.	

Prepare VNX

VNX5400 Unified Installation Guide provides instructions on assembly, racking, cabling, and powering the VNX. For 2,000 virtual desktops, refer to the *VNX5600 Unified Installation Guide* instead. There are no specific setup steps for this solution.

Set up the initial VNX configuration

After completing the initial VNX setup, configure key information about the existing environment so that the storage array can communicate. Configure the following items in accordance with your IT datacenter policies and existing infrastructure information:

- DNS
- NTP
- Storage network interfaces
- Storage network IP address
- CIFS services and Active Directory Domain membership

The reference documents listed in Table 30 provide more information on how to configure the VNX platform. Storage configuration guidelines on page 73 provides more information on the disk layout.

Provision core data storage

Core storage layout

Figure 23 shows the target storage layout for both FC and NFS variants for 500 virtual desktops. Figure 24 shows the target storage layout for both FC and NFS variants for 1,000 virtual desktops. Figure 27 shows the target storage layout for both FC and NFS variants for 2,000 virtual desktops.

Provision storage for VMFS datastores (block only)

Complete the following steps in EMC Unisphere to configure FC LUNs on VNX for storing virtual desktops:

1. Create a block-based RAID 5 storage pool that consists of 10 (for 500 virtual desktops) or 15 (for 1,000 virtual desktops) or 30 (for 2,000 virtual desktops) 300 GB SAS drives. Enable FAST Cache for the storage pool.
   a. Login to EMC Unisphere.
   b. Choose the array that will be used in this solution.
   c. Go to Storage > Storage Configuration > Storage Pools.
   d. Go to the Pools tab.
   e. Click Create.

Note: Create Hot Spare disks at this point. The *EMC VNX5400 Unified Installation Guide* and the *EMC VNX5600 Unified Installation Guide* provide additional information.

2. Configure 1 LUN of 50 GB and 4 LUNs of 485 GB (for 500 virtual desktops), 2 LUNs of 50 GB and 8 LUNs of 360 GB (for 1,000 virtual desktops), or 2 LUNs of 50 GB and 16 LUNs of 360 GB (for 2,000 virtual

desktops). Configure each LUN from the pool to present to the vSphere servers as four VMFS datastores.

a. Go to Storage > LUNs.

b. In the dialog box, click Create.

c. Select the Pool created in Step 1.

You will provision LUNs after this operation.

3. Configure a storage group to allow vSphere servers to access the newly created LUNs.

a. Go to Hosts > Storage Groups.

b. Create a new storage group.

c. Select LUNs and ESXi hosts to add to the storage group.

Provision storage for NFS datastores (file only)

Complete the following steps in EMC Unisphere to configure NFS file systems on VNX to store virtual desktops:

1. Create a block-based RAID 5 storage pool that consists of 10 (for 500 virtual desktops), 15 (for 1,000 virtual desktops), or 30 (for 2,000 virtual desktops) 300 GB SAS drives. Enable FAST Cache for the storage pool.

a. Login to EMC Unisphere.

b. Select the array that will be used in this solution.

c. Go to Storage > Storage Configuration > Storage Pools.

d. Go to the **Pools** tab.

e. Click **Create**.

Note: Create Hot Spare disks at this point. Refer to the *EMC VNX5400 Unified Installation Guide* for additional information.

---

2. Configure 10 LUNs of 200 GB (for 500 virtual desktops), 300 GB (for 1,000 virtual desktops), or 600 GB (for 2,000 virtual desktops) each from the pool to present to the Data Mover as dvols of a system-defined NAS pool.

a. Select **Storage** > **LUNs**.

b. Click **Create**.

c. Choose the pool created in Step 1, specify the corresponding user capacity, and indicate the number of LUNs you want to create.

d. Select **Hosts** > **Storage Groups**.

e. Select **file storage**.

f. Under Available LUNs, click **Connect LUNs**.

g. Choose the 10 LUNs you just created.

They appear in the Selected LUNs pane.

h. Select **A new storage pool for file is ready or manually rescan**.

i. Click **Storage** > **Storage Pool for File** > **Rescan Storage System** to create multiple file systems.

Note: EMC Performance Engineering best practice recommends that you create approximately 1 LUN for every 4 drives in the storage pool and that you create LUNs in even multiples of 10. Refer to *EMC VNX Unified Best Practices For Performance Applied Best Practices Guide*.

---

3. Configure four file systems of 485 GB each and one file system of 50 GB (for 500 virtual desktops), eight file systems of 360 GB each and two file systems of 50 GB each (for 1,000 virtual desktops), or 16 file systems of 365 GB each and two file systems of 50 GB each (for 2,000 virtual desktops). Configure these from the NAS pool to present to the vSphere servers as four NFS datastores.

a. Go to Storage > Storage Configuration > File Systems.

b. In the dialog box, click Create.

c. Select Create from Storage Pool.

d. Type a value in Storage Capacity and accept the default values for all other parameters.

4. Export the file systems using NFS, and give root access to vSphere servers.

5. In Unisphere:

a. Select **Settings** > **Data Mover Parameters** to make changes to the Data Mover configuration.

b. From the **Set Parameters** list, select **All Parameters**, as shown in Figure 42.

Figure 42. View all Data Mover parameters

c. Scroll down to the nthreads parameter as shown in Figure 43.

d. Right-click and select **Properties** to update the setting.

Note: The default number of threads serving NFS requests is 384 per data mover on VNX. Because more than 384 desktop connections are required in this solution, increase the number of active NFS threads to a maximum of 512 (for 500 virtual desktops) or 1,024 (for 1,000 virtual desktops) or 2,048 (for 2,000 virtual desktops) on each data mover.

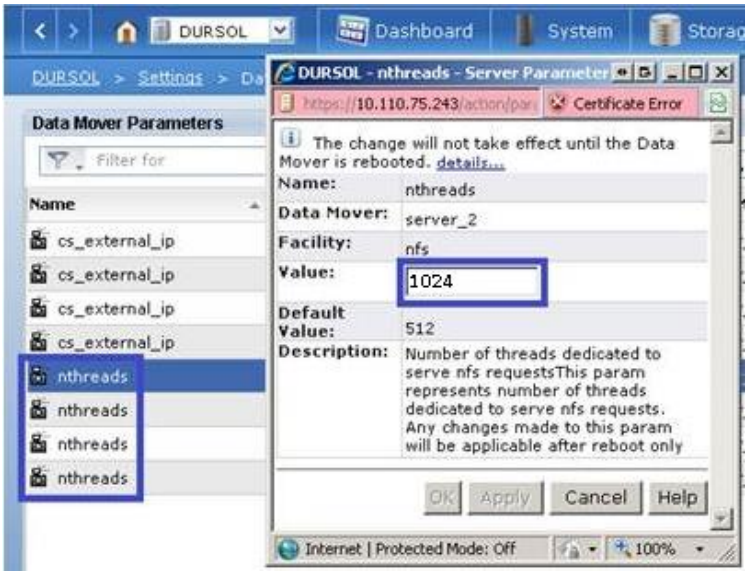

Figure 43. Set nthread parameter

Fast Cache configuration

To configure FAST Cache on the storage pool(s) for this solution complete the following steps.

1. Configure flash drives as FAST Cache:
   a. To create FAST Cache, from the Unisphere dashboard, click Properties or in the left pane of the Unisphere window, select Manage Cache.
   b. In the Storage System Properties dialog box, shown in Figure 44, select FAST Cache to view FAST Cache information.

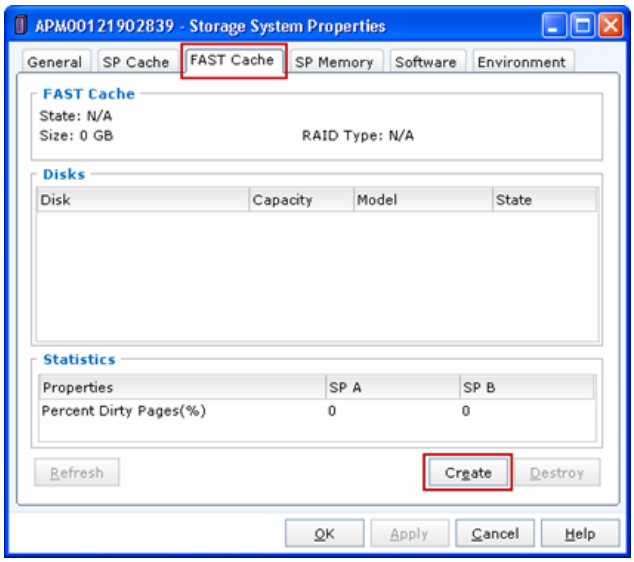

Figure 44. Storage System Properties dialog box

2. Click **Create** to open the Create FAST Cache dialog box, shown in Figure 45.

The RAID Type field is displayed as RAID 1 when the FAST Cache has been created. The number of flash drives can also be chosen in the screen. The bottom portion of the screen shows the flash drives that will be used for creating FAST Cache. You can choose the drives manually by selecting the **Manual** option.

3. See Storage configuration guidelines to determine the number of flash drives to use in this solution.

Note: If a sufficient number of flash drives is not available, an error message appears and FAST Cache cannot be created.

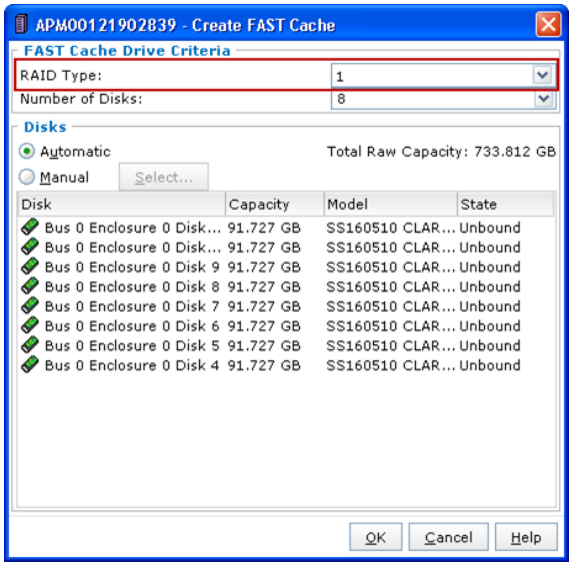
APM00121902839 - Create FAST Cache
FAST Cache Drive Criteria
RAID Type:
1
Number of Disks:
8
Disks
Automatic
Total Raw Capacity: 733.812 GB
Manual
Select...
Disk
Capacity
Model
State
Bus 0 Enclosure 0 Disk... 91.727 GB SS160510 CLAR... Unbound
Bus 0 Enclosure 0 Disk... 91.727 GB SS160510 CLAR... Unbound
Bus 0 Enclosure 0 Disk 9 91.727 GB SS160510 CLAR... Unbound
Bus 0 Enclosure 0 Disk 8 91.727 GB SS160510 CLAR... Unbound
Bus 0 Enclosure 0 Disk 7 91.727 GB SS160510 CLAR... Unbound
Bus 0 Enclosure 0 Disk 6 91.727 GB SS160510 CLAR... Unbound
Bus 0 Enclosure 0 Disk 5 91.727 GB SS160510 CLAR... Unbound
Bus 0 Enclosure 0 Disk 4 91.727 GB SS160510 CLAR... Unbound
OK
Cancel
Help

Figure 45. Create FAST Cache dialog box

4. Enable FAST Cache on the storage pool.

If a LUN is created in a storage pool, you can only configure FAST Cache for that LUN at the storage pool level. In other words, all the LUNs created in the storage pool will have FAST Cache enabled or disabled. You can configure them under the **Advanced** tab in the **Create Storage Pool** dialog box shown in Figure 46. After installing FAST Cache into the VNX array, it is enabled as a default setting when you create the storage pool.

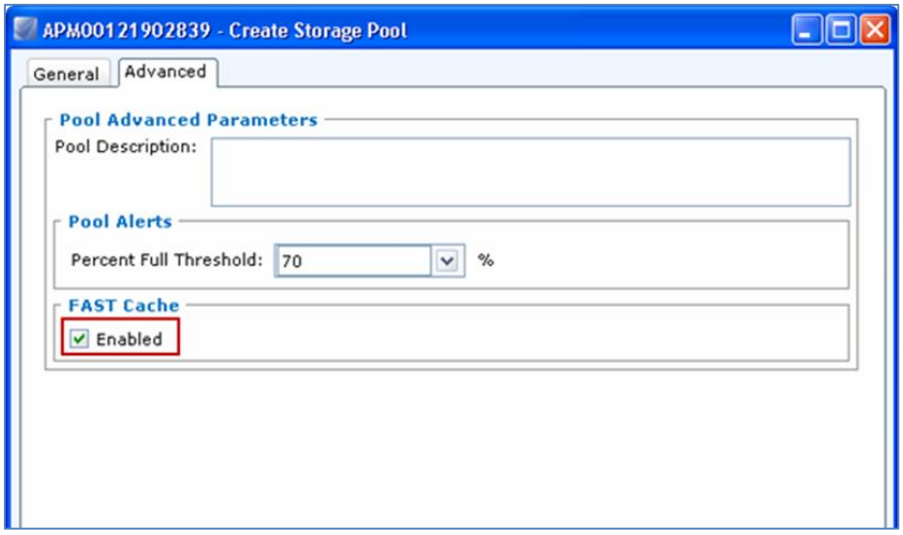

Figure 46. Create Storage Pool dialog box Advanced tab

If the storage pool has already been created, you can use the **Advanced** tab in the **Storage Pool Properties** dialog box to configure FAST Cache as shown in Figure 47.

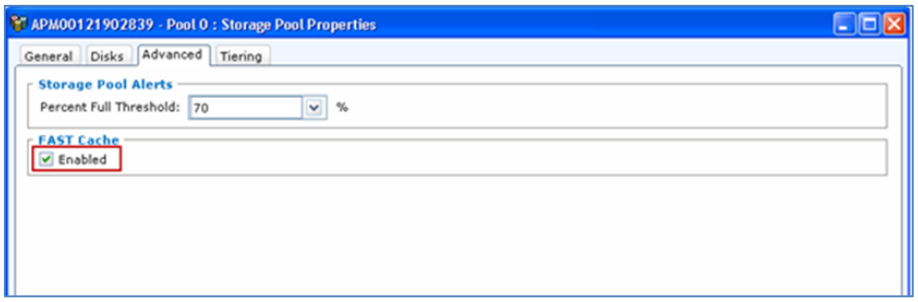

Figure 47. Storage Pool Properties dialog box Advanced tab

Note: The FAST Cache feature on the VNX series array does not cause an instantaneous performance improvement. The system must collect data about access patterns and promote frequently used information into the cache. This process can take a few hours during which the performance of the array steadily improves.

Provision optional storage for user data

If storage required for user data (i.e. roaming user profiles or View Persona Management repositories, and home directories) does not exist in the production environment already and the optional user data disk pack has been purchased, complete the following steps in Unisphere to configure two CIFS file systems on VNX:

1. Create a block-based RAID 6 storage pool that consists of 8 (for 500 virtual desktops), 16 (for 1,000 virtual desktops), or 32 (for 2,000 virtual desktops) 2 TB NL-SAS drives.

Figure 26 on page 86 depicts the target user data storage layout for 1,000 virtual desktops. Figure 28 on page 89 depicts the target user data storage layout for 2,000 virtual desktops.

2. Provision ten 1 TB (for 500 virtual desktops), 1.5 TB (for 1,000 virtual desktops), or 3 TB (for 2,000 virtual desktops) LUNs each from the pool to present to Data Mover as dvols that belong to a system-defined NAS pool.

3. Provision four file systems from the NAS pool to be exported as CIFS shares on a CIFS server.

FAST VP Configuration (optional)

Optionally, you can configure FAST VP to automate data movement between storage tiers. You can configure FAST VP in two ways:

- To View and manage FAST VP at the pool level, select a storage pool and click **Properties** to open the **Storage Pool Properties** dialog box.

Figure 48 shows the tiering information for a specific FAST VP enabled pool.

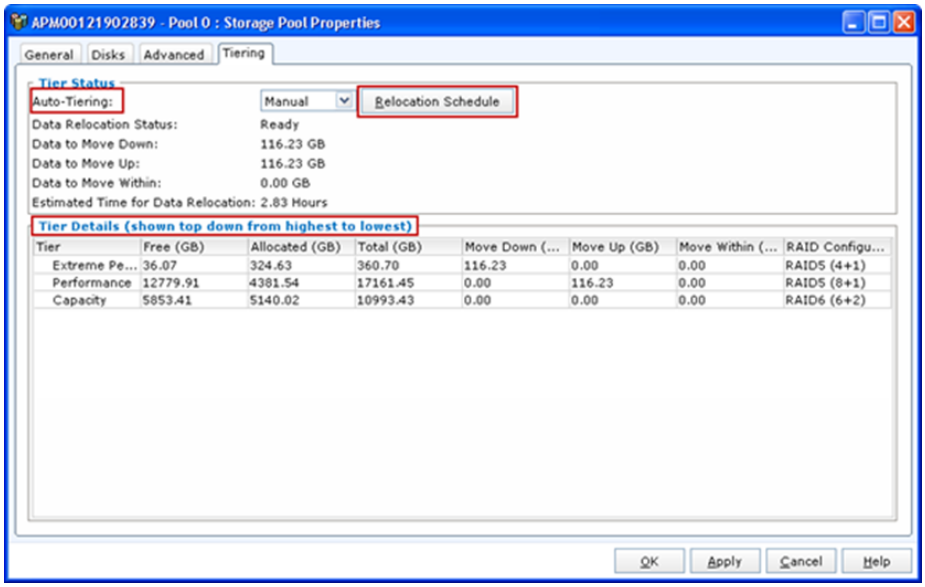
APM00121902839 - Pool 0 : Storage Pool Properties
General
Disks
Advanced
Tiering
Tier Status
Auto-Tiering:
Manual
Relocation Schedule
Data Relocation Status: Ready
Data to Move Down: 116.23 GB
Data to Move Up: 116.23 GB
Data to Move Within: 0.00 GB
Estimated Time for Data Relocation: 2.83 Hours
Tier Details (shown top down from highest to lowest)
Tier | Free (GB) | Allocated (GB) | Total (GB) | Move Down (... | Move Up (GB) | Move Within (... | RAID Configu...
Extreme Pe... | 36.07 | 324.63 | 360.70 | 116.23 | 0.00 | 0.00 | RAID5 (4+1)
Performance | 12779.91 | 4381.54 | 17161.45 | 0.00 | 116.23 | 0.00 | RAID5 (8+1)
Capacity | 5853.41 | 5140.02 | 10993.43 | 0.00 | 0.00 | 0.00 | RAID6 (6+2)
OK
Apply
Cancel
Help

Figure 48. Storage Pool Properties dialog box

- The Tier Status section shows FAST VP relocation information specific to the pool selected. Scheduled relocation can be selected at the pool level from the Auto-Tiering menu. This can be set to either Automatic or Manual. In the Tier Details section, you can see the exact distribution of the data. Click Relocation Schedule to open the Manage Auto-Tiering window shown in Figure 49.

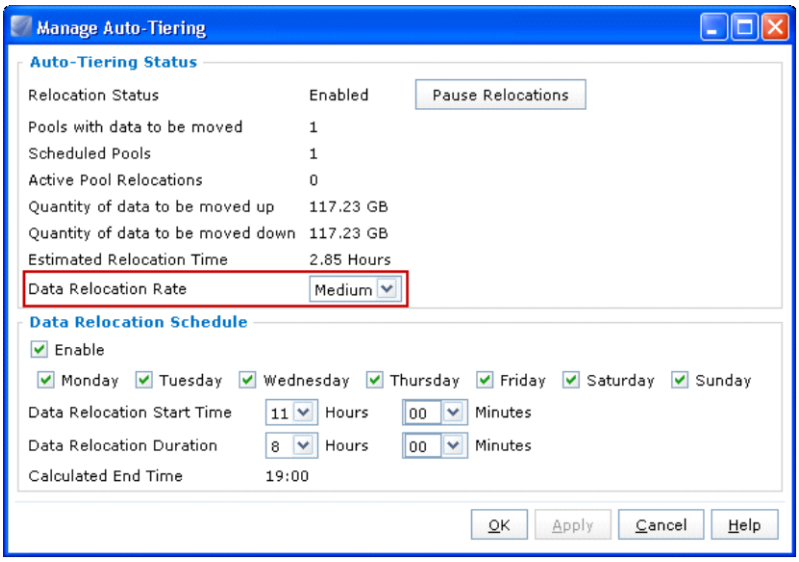

Figure 49. Manage Auto-Tiering Window

From this status window, you can control the Data Relocation Rate. The default rate is set to Medium to avoid significantly affecting host I/O.

Note: FASTVP is a completely automated tool and you can schedule relocations to occur automatically. EMC recommends that relocations be scheduled during off-hours to minimize any potential performance impact.

- Configure FAST VP at the LUN level. Some FAST VP properties are managed at the LUN level. Click Properties of a specific LUN. In this dialog box, click the Tiering tab to view tiering information for this single LUN as shown in Figure 50.

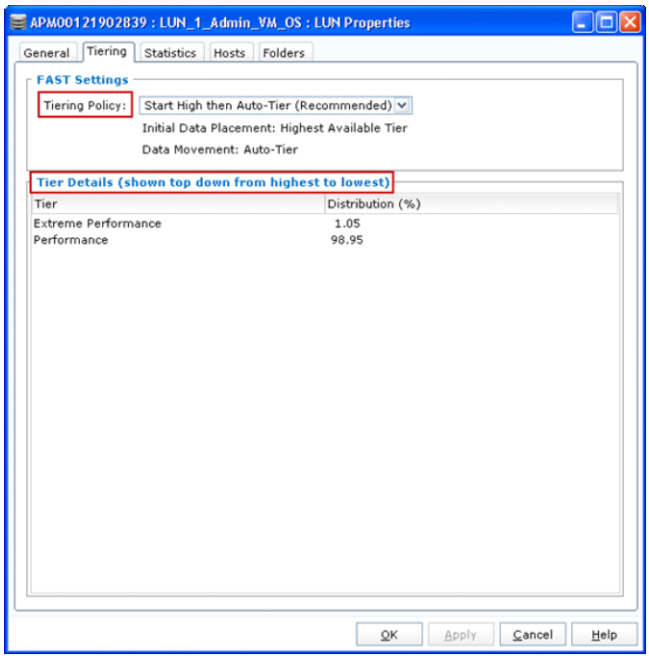

Figure 50. LUN Properties window

The Tier Details section displays the current distribution of slices within the LUN. Tiering policy can be selected at the LUN level from the Tiering Policy list.

Provision optional storage for infrastructure virtual machines

If the storage required for infrastructure virtual machines (that is, SQL server, domain controller, vCenter server, and/or VMware Horizon View Connection Servers) does not exist in the production environment already and you have purchased the optional user data disk pack, configure a NFS file system on VNX to be used as the NFS datastore in which the infrastructure virtual machines reside. Repeat the configuration steps that are shown in Provision storage for NFS datastores (file only) to provision the optional storage, while taking into account the smaller number of drives.

Install and configure vSphere hosts

Overview

This section provides information about the installation and configuration of vSphere hosts and infrastructure servers required to support the architecture. Table 31 describes the tasks to be completed.

Table 31. Tasks for server installation

Task	Description	Reference
Install vSphere	Install the vSphere hypervisor on the physical servers deployed for the solution.	vSphere Installation and Setup Guide
Configure vSphere networking	Configure vSphere networking including NIC trunking, VMkernel ports, and virtual machine port groups and Jumbo Frames.	vSphere Networking
Add vSphere hosts to VNX storage groups (FC variant)	Use the Unisphere console to add the vSphere hosts to the storage groups created in Prepare and configure the storage .	
Connect VMware datastores	Connect the VMware datastores to the vSphere hosts deployed for the solution.	vSphere Storage Guide

Install vSphere

Upon initial power up of the servers being used for vSphere, confirm or enable the hardware-assisted CPU virtualization and the hardware-assisted MMU virtualization setting in the server's BIOS. If the servers are equipped with a RAID controller, we recommend that you configure mirroring on the local disks.

Start up the vSphere installation media and install the hypervisor on each of the servers. vSphere hostnames, IP addresses, and a root password are required for installation. Appendix B provides appropriate values.

Configure vSphere NetworkingConniknfnevjrevnervlkjnvrelgv

During the installation of VMware vSphere, a standard virtual switch (vSwitch) is created. By default, vSphere chooses only one physical NIC as a vSwitch uplink. To maintain redundancy and bandwidth requirements, configure an additional NIC, either by using the vSphere console or by connecting to the vSphere host from the vSphere Client.

Each VMware vSphere server should have multiple interface cards for each virtual network to ensure redundancy and provide for the use of network load balancing, link aggregation, and network adapter failover.

VMware vSphere networking configuration, including load balancing, link aggregation, and failover options, is described in *vSphere Networking*. Refer to the list of documents in Appendix C of this document for more information. Choose the appropriate load-balancing option based on what is supported by the network infrastructure.

Create VMkernel ports as required, based on the infrastructure configuration:

- VMkernel port for NFS traffic (NFS variant only)
- VMkernel port for VMware vMotion
- Virtual desktop port groups (used by the virtual desktops to communicate on the network)

vSphere Networking describes the procedure for configuring these settings. Refer to the list of documents in Appendix C of this document for more information.

Jumbo frames

A jumbo frame is an Ethernet frame with a "payload" greater than 1,500 bytes and up to 9,000 bytes. This is also known as the maximum transmission unit (MTU). The generally accepted maximum size for a jumbo frame is 9,000 bytes. Processing overhead is proportional to the number of frames. Therefore, enabling jumbo frames reduces processing overhead by reducing the number of frames to be sent. This increases the network throughput. Jumbo frames must be enabled end-to-end. This includes the network switches, vSphere servers, and VNX SPs. EMC recommends enabling jumbo frames on the networks and interfaces used for carrying NFS traffic.

Jumbo frames can be enabled on the vSphere server into two different levels. If all the portals on the vSwitch need to be enabled for jumbo frames, this can be achieved by selecting Properties of the vSwitch and editing the MTU settings from the vCenter. If specific VMkernel ports are to be jumbo frame-enabled, edit the VMkernel port under Network Properties from the vCenter.

To enable jumbo frames on the VNX:

1. Navigate to Unisphere >Settings > Network > Settings for File.
2. Select the appropriate network interface under the **Interfaces** tab.
3. Select **Properties**.
4. Set the MTU size to **9,000**.
5. Click **OK** to apply the changes.

Jumbo frames might also need to be enabled on each network switch. Consult your switch configuration guide for instructions.

Connect VMware datastores

Connect the data stores configured in Prepare and configure the storage to the appropriate vSphere servers. These include the datastores configured for:

- Virtual desktop storage
- Infrastructure virtual machine storage (if required)
- SQL Server storage (if required)

vSphere Storage Guide provides instructions on how to connect the VMware datastores to the vSphere host. Refer to the list of documents in Appendix C of this document for more information.

The vSphere EMC PowerPath VE (FC variant) and NFS VAAI (NFS variant) plug-ins must be installed after VMware Virtual Center has been deployed as described in

VMware vCenter Server Deployment.

Plan virtual machine memory allocations

Server capacity is required for two purposes in the solution:

- To support the new virtualized desktop infrastructure
- To support the required infrastructure services such as authentication/authorization, DNS, and database

For information on minimum infrastructure services hosting requirements, refer to Table 6. If existing infrastructure services meet the requirements, the hardware listed for infrastructure services is not required.

Memory configuration

Proper sizing and configuration of the solution requires that you are careful when configuring server memory. The following section provides general guidance on memory allocation for the virtual machines and factors in vSphere overhead and the virtual machine configuration.

ESX/ESXi memory management

Memory virtualization techniques allow the vSphere hypervisor to abstract physical host resources, such as memory, to provide resource isolation across multiple virtual machines, while avoiding resource exhaustion. In cases where advanced processors (such as Intel processors with EPT support) are deployed, this abstraction takes place within the CPU. Otherwise, this process occurs within the hypervisor itself using a feature known as shadow page tables.

vSphere employs the following memory management techniques:

Allocation of memory resources greater than those physically available to the virtual machine is known as *memory over-commitment*.

Identical memory pages that are shared across virtual machines are merged using a feature known as *transparent page sharing*. Duplicate pages are returned to the host's free memory pool for reuse.

Memory compression is a process in which ESXi stores pages, which would otherwise be swapped out to disk through host swapping, in a compression cache located in the main memory.

Host resource exhaustion can be relieved using a process known as *memory ballooning*, which allocates free pages from the virtual machine to the host for reuse.

Hypervisor swapping causes the host to force arbitrary virtual machine pages out to disk.

You can find additional information on VMware's website at www.vmware.com/files/pdf/mem_mgmt_perf_vsphere5.pdf.

Virtual machine memory concepts

Figure 51 shows the memory settings parameters in the virtual machine, including:

- **Configured memory—**Physical memory allocated to the virtual machine at the time of creation
- **Reserved memory—**Memory that is guaranteed to the virtual machine
- **Touched memory—**Memory that is active or in use by the virtual machine
- **Swappable—**Memory that can be de-allocated from the virtual machine if the host is under memory pressure from other virtual machines using ballooning, compression, or swapping

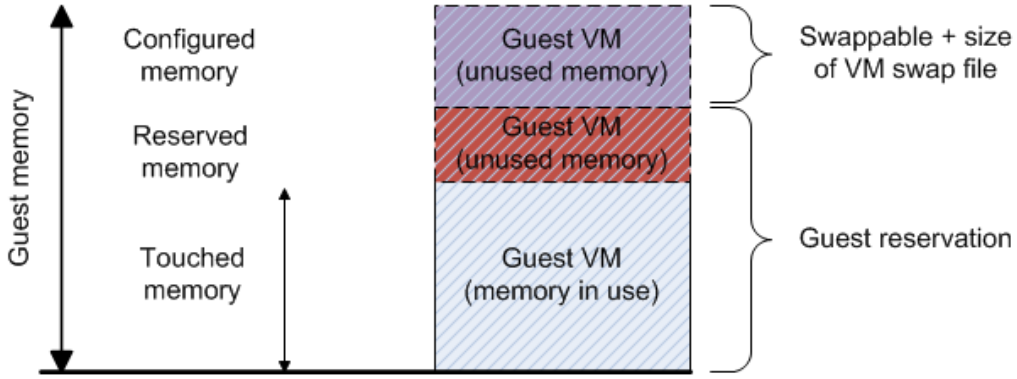

Figure 51. Virtual machine memory settings

We recommend that you follow these best practices:

- Do not disable the default memory reclamation techniques. These lightweight processes enable flexibility with minimal impact to workloads.
- Intelligently size memory allocation for virtual machines. Over-allocation wastes resources, while under-allocation causes performance impacts that can affect other virtual machines sharing resources. Over-committing can lead to resource exhaustion if the hypervisor cannot procure memory resources. In severe cases when hypervisor swapping is encountered, virtual machine performance will likely be adversely affected. Having performance baselines of your virtual machine workloads assists in this process.

Install and configure SQL Server database

Overview

This section and Table 32 describe how to set up and configure a Microsoft SQL Server database for the solution. When the steps in this section have been completed, Microsoft SQL Server will be on a virtual machine with the databases required by VMware vCenter, VMware Update Manager, VMware Horizon View, and VMware View Composer configured for use.

Table 32. Tasks for SQL Server database setup

Task	Description	Reference
Create a virtual machine for Microsoft SQL Server	Create a virtual machine to host SQL Server. Verify that the virtual server meets the hardware and software requirements.	msdn.microsoft.com
Install Microsoft Windows on the virtual machine	Install Microsoft Windows Server 2008 R2 Standard Edition on the virtual machine created to host SQL Server.	technet.microsoft.com
Install Microsoft SQL Server	Install Microsoft SQL Server on the virtual machine designated for that purpose.	technet.microsoft.com
Configure database for VMware vCenter	Create the database required for the vCenter Server on the appropriate datastore.	Preparing vCenter Server Databases
Configure database for VMware Update Manager	Create the database required for Update Manager on the appropriate datastore.	Preparing the Update Manager Database
Configure database for VMware Horizon View Composer	Create the database required for View Composer on the appropriate datastore.	VMware Horizon View 5.3 Installation
Configure database for VMware Horizon View Manager	Create the database required for VMware Horizon View Manager event logs on the appropriate datastore.	VMware Horizon View 5.3 Installation

Task	Description	Reference
Configure the VMware Horizon View and View Composer database permissions	Configure the database server with appropriate permissions for the VMware Horizon View and VMware Horizon View Composer databases.	VMware Horizon View 5.3 Installation
Configure VMware vCenter database permissions	Configure the database server with appropriate permissions for the VMware vCenter.	Preparing vCenter Server Databases
Configure VMware Update Manager database permissions	Configure the database server with appropriate permissions for the VMware Update Manager.	Preparing the Update Manager Database

Create a virtual machine for Microsoft SQL Server

Create the virtual machine with enough computing resources on one of the Windows servers designated for infrastructure virtual machines, and use the datastore designated for the shared infrastructure.

Note: The customer environment may already contain an SQL Server that is designated for this role. In that case, refer to Configure the database for VMware vCenter on page 169.

Install Microsoft Windows on the virtual machine

The SQL Server service must run on Microsoft Windows. Install Windows on the virtual machine by selecting the appropriate network, time, and authentication settings.

Install SQL Server

Install SQL Server on the virtual machine from the SQL Server installation media. The Microsoft TechNet website provides information on how to install SQL Server.

One of the installable components in the SQL Server installer is the SQL Server Management Studio (SSMS). You can install this component on the SQL server directly as well as on an administrator's console. SSMS must be installed on at least one system.

You have the option to store data files in locations other than the default path. To change the default path, right-click on the server object in SSMS and select **Database Properties**. This action opens a properties dialog box from which you can change the default data and log directories for new databases created on the server.

Note: For high availability, an SQL Server can be installed in a Microsoft Failover Clustering or on a virtual machine protected by VMHA clustering. Do not combine these technologies.

Configure the database for VMware vCenter

To use VMware vCenter in this solution, create a database for the service to use. The requirements and steps to configure the vCenter Server database correctly are covered in *Preparing vCenter Server Databases*. Refer to the list of documents in Appendix C of this document for more information.

Note: Do not use the Microsoft SQL Server Express-based database option for this solution.

It is a best practice to create individual login accounts for each service accessing a database on SQL Server.

Configure database for VMware Update Manager

To use VMware Update Manager in this solution, create a database for the service to use. The requirements and steps to configure the Update Manager database correctly are covered in *Preparing the Update Manager Database*. Refer to the list of documents in Appendix C of this document for more information. It is a best practice to create individual login accounts for each service accessing a database on SQL Server. Consult your database administrator for your organization's policy.

Configure database for VMware View Composer

To use VMware View Composer in this solution, create a database for the service to use. The requirements and steps to configure the Update Manager database correctly are covered in *VMware Horizon View 5.3 Installation*. Refer to the list of documents in Appendix C of this document for more information. It is a best practice to create individual login accounts for each service accessing a database on SQL Server. Consult your database administrator for your organization's policy.

Configure database for VMware Horizon View Manager

To retain VMware Horizon View event logs, create a database for the VMware Horizon View Manager to use. *VMware Horizon View 5.3 Installation* provides the requirements and steps to configure the VMware Horizon View event database correctly. Refer to the list of documents in Appendix C of this document for more information. It is a best practice to create individual login accounts for each service accessing a database on SQL Server. Consult your database administrator for your organization's policy.

Configure the VMware Horizon View and View Composer database permissions

At this point, your database administrator must create user accounts that will be used for the View Manager and View Composer databases and assign the appropriate permissions. It is a best practice to create individual login accounts for each service accessing a database on SQL Server. Consult your database administrator for your organization's policy.

VMware vCenter Server Deployment

Overview

This section provides information on how to configure VMware vCenter. Table 33 describes the tasks to be completed.

Table 33. Tasks for vCenter configuration

Task	Description	Reference
Create the vCenter host virtual machine	Create a virtual machine for the VMware vCenter Server.	vSphere Virtual Machine Administration
Install vCenter guest OS	Install Windows Server 2008 R2 Standard Edition on the vCenter host virtual machine.	
Update the virtual machine	Install VMware Tools, enable hardware acceleration, and allow remote console access.	vSphere Virtual Machine Administration
Create vCenter ODBC connections	Create the 64-bit vCenter and 32-bit vCenter Update Manager ODBC connections.	• vSphere Installation and Setup • Installing and Administering VMware vSphere Update Manager
Install vCenter Server	Install vCenter Server software.	vSphere Installation and Setup
Install vCenter Update Manager	Install vCenter Update Manager software.	Installing and Administering VMware vSphere Update Manager
Create a virtual datacenter	Create a virtual datacenter.	vCenter Server and Host Management
Apply vSphere license keys	Type the vSphere license keys in the vCenter licensing menu.	vSphere Installation and Setup
Add vSphere Hosts	Connect the vCenter to the vSphere hosts.	vCenter Server and Host Management
Configure vSphere clustering	Create a vSphere cluster and move the vSphere hosts into it.	vSphere Resource Management
Perform array vSphere host discovery	Perform vSphere host discovery within the Unisphere console.	Using EMC VNX Storage with VMware vSphere–TechBook

Task	Description	Reference
Install the vCenter Update Manager plug-in	Install the vCenter Update Manager plug-in on the administration console.	Installing and Administering VMware vSphere Update Manager
vStorage APIs for Array Integration (VAAI) Plug-in	Using VMware Update Manager, deploy the vStorage APIs for Array Integration (VAAI) plug-in to all vSphere hosts.	• EMC VNX VAAI NFS Plug-in– Installation HOWTO video available on www.youtube.com • vSphere Storage APIs for Array Integration (VAAI) Plug-in • Installing and Administering VMware vSphere Update Manager
Deploy PowerPath/VE (FC Variant)	Use VMware Update Manager to deploy PowerPath/VE plug-in to all vSphere hosts.	PowerPath/VE for VMware vSphere Installation and Administration Guide
Install the EMC VNX UEM CLI	Install the EMC VNX UEM CLI on the administration console.	*EMC VSI for VMware vSphere: Unified Storage Management— Product Guide*
Install the EMC VSI plug-in	Install the EMC Virtual Storage Integration plug-in on the administration console.	*EMC VSI for VMware vSphere: Unified Storage Management— Product Guide*
Install the EMC PowerPath Viewer (FC Variant)	Install the EMC PowerPath Viewer on the administration console.	PowerPath Viewer Installation and Administration Guide

Create the vCenter host virtual machine

If the VMware vCenter Server is to be deployed as a virtual machine on a vSphere server installed as part of this solution, connect directly to an Infrastructure vSphere server using the vSphere Client. Create a virtual machine on the vSphere server with the guest OS configuration using the infrastructure server datastore presented from the storage array. The memory and processor requirements for the vCenter server are dependent on the number of vSphere hosts and virtual machines being managed. The requirements are outlined in the *vSphere Installation and Setup Guide*. Refer to the list of documents in Appendix C of this document for more information.

Install vCenter guest OS

Install the guest OS on the vCenter host virtual machine. VMware recommends using Windows Server 2008 R2 Standard Edition. Refer to the list of documents in Appendix C of this document for more information.

Create vCenter ODBC connections

Before installing the vCenter Server and vCenter Update Manager, create the ODBC connections required for database communication. These ODBC connections will use SQL Server authentication for database authentication. The section entitled Configure the database for VMware vCenter on page 169 provides SQL login information.

Refer to the list of documents in Appendix C of this document for more information.

Install vCenter Server

Install vCenter by using the VMware VIMSetup installation media. Use the customer-provided username, organization, and vCenter license key when installing vCenter.

Apply vSphere license keys

To perform license maintenance, log in to the vCenter Server and select the **Administration** > **Licensing** menu from the vSphere Client. Use the vCenter License console to enter the license keys for the vSphere hosts. After this, they can be applied to the vSphere hosts as they are imported into vCenter.

vStorage APIs for Array Integration (VAAI) Plug-in

The vStorage APIs for Array Integration (VAAI) plug-in enables support for the vSphere NFS primitives. These primitives reduce the load on the hypervisor from specific storage-related tasks to free resources for other operations. Additional information about the VAAI for NFS plug-in is available in the plug-in download *vSphere Storage APIs for Array Integration (VAAI) Plug-in*. Refer to the list of documents in Appendix C of this document for more information.

The VAAI for NFS plug-in is installed using vSphere Update Manager. Refer to the process for distributing the plug-in demonstrated in the EMC VNX VAAI NFS plug-in – installation HOWTO video available on the www.youtube.com website. To enable the plug-in after installation, restart the vSphere server.

Deploy PowerPath/VE (FC variant)

EMC PowerPath is host-based software that provides automated data path management and load-balancing capabilities for heterogeneous server, network, and storage deployed in physical and virtual environments. PowerPath uses multiple I/O data paths to share the workload, and automated load balancing to ensure the efficient use of data paths.

The PowerPath/VE plug-in is installed using the vSphere Update Manager. *PowerPath/VE for VMware vSphere Installation and Administration Guide* describes the process to distribute the plug-in and apply the required licenses. To enable the plug-in after installation, restart the vSphere server.

Install the EMC VSI plug-in

The VNX storage system can be integrated with VMware vCenter using the EMC Virtual Storage Integrator (VSI) for VMware vSphere Unified Storage Management plug-in. This provides administrators the ability to manage VNX storage tasks from within the vSphere Client.

After the plug-in is installed on the vSphere console, administrators can use vCenter to:

- Create datastores on VNX and mount them on vSphere servers.
- Extend datastores.
- Create FAST or full clones of virtual machines.

Set Up VMware View Connection Server

Overview

This section provides information on how to set up and configure VMware View Connection Servers for the solution. For a new installation of VMware Horizon View, VMware recommends that you complete the following tasks in the order shown in Table 34.

Table 34. Tasks for VMware Horizon View Connection Server setup

Task	Description	Reference
Create virtual machines for VMware View Connection Servers	Create two virtual machines in vSphere Client. These virtual machines will be used as VMware View Connection Servers.	VMware Horizon View 5.3 Installation
Install guest OS for VMware View Connection Servers	Install Windows Server 2008 R2 guest OS.	
Install VMware View Connection Server	Install VMware View Connection Server software on one of the previously prepared virtual machines.	VMware Horizon View 5.3 Installation
Enter the View license key	Enter the View license key in the View Manager web console.	
Configure the View event log database connection	Configure the View event log database settings using the appropriate database information and login credentials.	
Add a replica View Connection Server	Install VMware View Connection Server software on the second server.	
Configure the View Composer ODBC connection	On either the vCenter Server or a dedicated Windows Server 2008 R2 server, configure an ODBC connection for the previously configured View Composer database.	
Install View Composer	Install VMware View Composer on the server identified in the previous step.	
Connect VMware Horizon View to vCenter and View Composer	Use the View Manager web interfaces to connect View to the vCenter server and View Composer.	VMware Horizon View 5.3 Administration
Prepare a master virtual machine	Create a master virtual machine as the base image for the virtual desktops.	
Configure View Persona Management Group Policies	Configure AD Group Policies to enable View Persona Management.	
Configure View PCoIP Group Policies	Configure AD Group Policies for PCoIP protocol settings.	

Install the VMware Horizon View Connection Server	Install the View Connection Server software using the instructions from *VMware Horizon View 5.3 Installation*. Select **Standard** when prompted for the View Connection Server type.
Configure the View Event Log Database connection	Configure the VMware Horizon View event log database connection using the database server name, database name, and database log in credentials. Review the *VMware Horizon View 5.3 Installation* guide for specific instructions on how to configure the event log.
Add a second View Connection Server	Repeat the View Connection Server installation process on the second target virtual machine. When prompted for the connection server type, specify **Replica**, and then provide the VMware Horizon View administrator credentials to replicate the View configuration data from the first View Connection Server.
Configure the View Composer ODBC connection	On the server that will host the View Composer service, create an ODBC connection for the previously configured View Composer database. Review the *VMware Horizon View 5.3 Installation* guide for specific instructions on how to configure the ODBC connection.
Install View Composer	On the server that will host the View Composer service, install the View Composer software. Specify the previously configured ODBC connection when prompted during the installation process. Review the *VMware Horizon View 5.3 Installation Guide* for specific instructions on how to configure the ODBC connection.
Link VMware Horizon View to vCenter and View Composer	Using the VMware Horizon View Manager web console, create the connection between View and both the vCenter Server and the View Composer. Review the *VMware Horizon View 5.3 Administration Guide* for specific instructions on how to create the connections. When presented with the option, enable vSphere host caching (also known as View Storage Accelerator or Content Based Read Cache) and set the cache amount at 2 GB, the maximum amount supported. You can also enable the **Reclaim VM** disk space option. This feature is currently supported only with Windows 7 desktops. If you enable the Reclaim VM disk space option, you must specify a blackout period that controls when the operation should not process. The operation should not execute during periods of heavy use, so you should specify that those times are included in the blackout period. By default, space reclamation will only run when there is 1 GB of space or more to reclaim. You configure this value when implementing your desktop pools.
Prepare master virtual machine	Optimize the master virtual machine to avoid unnecessary background services generating extraneous I/O operations that adversely affect the overall performance of the storage array.

Complete the following steps to prepare the master virtual machine:

1. Using the VMware vSphere Web Client, create a virtual machine using the VMware version 9 hardware specification. You cannot create version 9 virtual machines with the software client; you must use the web client.
2. Install Windows 7 guest OS.
3. Install appropriate integration tools such as VMware Tools.
4. Optimize the OS settings by referring the following documents: *Deploying Microsoft Windows 7 Virtual Desktops with VMware Horizon View —Applied Best Practices white paper*, and the *VMware Horizon View Optimization Guide for Windows 7 white paper*.
5. Install third-party tools or applications, such as Microsoft Office, relevant to your environment.
6. Install the Avamar Desktop/Laptop Client (Refer to
7. Set Up EMC Avamar for details).
8. Install the VMware Horizon View agent.

Note: If the View Persona Management feature will be used, the Persona Management component of the VMware Horizon View agent should be installed at this time. Ensure that the Persona Management option is selected during the installation of the View agent.

Configure View Persona Management Group Policies

View Persona Management is enabled using Active Directory Group Policies that are applied to the Organizational Unit (OU) containing the virtual desktop computer accounts. The View Group Policy templates are located in the *\Program Files\VMware\VMware Horizon View\Server\extras\GroupPolicyFiles* directory on the View Connection Server.

Configure Folder Redirection Group Policies for Avamar

Folder redirection is enabled using Active Directory Group Policies that are applied to the Organizational Unit (OU) containing the virtual desktop user accounts. The Active Directory folder redirection is used (instead of View Persona Management folder redirection) to ensure that the folders maintain the naming consistencies required by the Avamar software. Refer to

Set Up EMC Avamar for details.

Configure View PCoIP Group Policies

You control View PCoIP protocol settings using Active Directory Group Policies that are applied to the OU containing the VMware View Connection Servers. The View Group Policy templates are located in the *\Program Files\VMware\VMware View\Server\extras\GroupPolicyFiles* directory on the View Connection Server. You should use the group policy template **pcoip.adm** to specify the following PCoIP protocol settings:

- Maximum Initial Image Quality value: 70
- Maximum Frame Rate value: 24
- Turn off Build-to-Lossless feature: Enabled

Note: Higher PCoIP session frame rates and image qualities can adversely affect server resources.

Set Up EMC Avamar

See Design and Implementation Guide: EMC Backup and Recovery Options for VSPEX End User Computing for VMware Horizon View available on EMC Online Support.

Set up VMware vShield Endpoint

Overview

This section provides information on how to set up and configure the components of vShield Endpoint.

Table 35 describes the tasks to be completed.

Table 35. Tasks required to install and configure vShield Endpoint

Task	Description	Reference
Verify desktop vShield Endpoint driver installation	Verify that the vShield Endpoint driver component of VMware Tools has been installed on the virtual desktop master image.	vShield Quick Start Guide
Deploy vShield Manager appliance	Deploy and configure the VMware vShield Manager appliance.	
Register the vShield Manager plug-in.	Register the vShield Manager plug-in with the vSphere Client.	
Apply vShield Endpoint licenses	Apply the vShield Endpoint license keys using the vCenter license utility.	
Install vSphere vShield Endpoint service	Install the vShield Endpoint service on the desktop vSphere hosts.	
Deploy an antivirus solution management server	Deploy and configure an antivirus solution management server.	*vShield Quick Start Guide* **Note:** vShield Endpoint partners provide the antivirus management server software and security virtual machines. Consult the vendor documentation for specific details concerning installation and configuration.
Deploy vSphere security virtual machines (SVMs)	Deploy and configure security virtual machines on each desktop vSphere host.	
Verify vShield Endpoint functionality	Verify functionality of vShield Endpoint components using the virtual desktop master image.	**Note:** Consult vendor documentation for specific details on how to verify vShield Endpoint integration and functionality.

Verify desktop vShield Endpoint driver installation

The vShield Endpoint driver is a subcomponent of the VMware Tools software package that is installed on the virtual desktop master image. The driver is installed using one of two methods:

- Select the **Complete** option during VMware Tools installation.
- Select the **Custom** option during VMware Tools installation. From the VMware Device Drivers list, select **VMCI Driver,** and then select **vShield Driver**.

To install the vShield Endpoint driver on a virtual machine that already has VMware Tools installed, simply initiate the VMware Tools installation and select the appropriate option.

Deploy vShield Manager appliance

The vShield Manager appliance is provided by VMware as an OVA file that is imported through the vShield client using the **File – Deploy OVF template** option. The vShield Manager appliance is preconfigured with all required components.

Install the vSphere vShield Endpoint service

The vSphere vShield Endpoint service must be installed on all vSphere virtual desktop hosts. The service is installed on the vSphere hosts by the vShield Manager appliance. The vShield Manager web console is used to initiate the vShield Endpoint service installation and to verify that the installation is successful.

Deploy an antivirus solution management server

The antivirus solution management server is used to manage the antivirus solution and is provided by vShield Endpoint partners. The management server and associated components are a required component of the vShield Endpoint platform.

Deploy vSphere security virtual machines

The vSphere security virtual machines are provided by the vShield Endpoint partners and are installed on each vSphere virtual desktop host. The security virtual machines perform security related operations for all virtual desktops that reside on their vSphere host. The security virtual machines and associated components are required components of the vShield Endpoint platform.

Verify vShield Endpoint functionality

Once all required components of the vShield Endpoint platform have been installed and configured, the functionality of the platform should be verified prior to the deployment of virtual desktops.

Using the documentation provided by the vShield Endpoint partner, verify the functionality of the vShield Endpoint platform with the virtual desktop master image.

Set Up VMware vCenter Operations Manager for View

Overview

This section provides information on how to set up and configure VMware vCOps for View.

Table 36 describes the tasks that must be completed.

Table 36. Tasks required to install and configure vCOps

Task	Description	Reference
Create vSphere IP Pool for vCOps	Create an IP pool with two available IPs.	Deployment and Configuration Guide – vCenter Operations Manager 5
Deploy vCOps vSphere Application Services (vApp)	Deploy and configure the vCOps vApp.	
Specify the vCenter server to monitor	From the vCenter Operations Manager main web interface, specify the name of the vCenter server that manages the virtual desktops.	
Assign the vCOps license	Apply the vCOps for View license keys using the vCenter license utility.	
Configure SNMP and SMTP settings	From the vCenter Operations Manager main web interface, configure any required SNMP or SMTP settings for monitoring purposes. Note: Optional.	
Update virtual desktop settings	Update virtual desktop firewall policies and services to support vCOps for View desktop-specific metrics gathering.	vCenter Operations Manager for View Integration Guide
Create the virtual machine for the vCOps for View Adapter server	Create a virtual machine in the vSphere Client. The virtual machine will be used as the vCOps for View Adapter server.	
Install guest OS for the vCOps for View Adapter server	Install Windows Server 2008 R2 guest OS.	
Install the vCOps for View Adapter software	Deploy and configure the VCOps for View Adapter software.	vCenter Operations Manager for

Task	Description	Reference
Import the vCOps for View PAK file	Import the vCenter Operations Manager for View Adapter PAK file using the vCOps main web interface.	View Integration Guide
Verify vCOps for View functionality	Verify functionality of vCOps for View using the virtual desktop master image.	

Create vSphere IP Pool for vCOps

vCOps requires two IP addresses for use by the vCOps analytics and user interface (UI) virtual machines. These IP addresses will be assigned to the servers automatically during the deployment of the vCOps vApp.

Deploy vCenter Operations Manager vApp

The vCOps vApp is provided by VMware as an OVA file that is imported through the vShield client using the **File – Deploy OVF template** menu option. The vApp must be deployed on a vSphere cluster with DRS enabled.

The specifications of the two virtual servers that comprise the vCOps vApp must be adjusted based on the number of virtual machines being monitored.

Specify the vCenter server to monitor

Access the vCOps web interface using the web address: http://*ip*/admin where *ip* is the IP address or fully qualified host name of the vCOps vApp. Log in using the default credentials: user name admin and password admin. Complete the vCOps **First Boot Wizard** to complete the initial vCOps configuration and specify the appropriate vCenter server to monitor.

Update virtual desktop settings

vCOps for View requires the ability to gather metric data directly from the virtual desktop. To enable this capability, the virtual desktop service and firewall settings must be adjusted either by using Windows group policies or by updating the configuration of the virtual desktop master image.

The following virtual desktop changes need made to support vCOps for View:

- Add the following programs to the Windows 7 firewall **allow** list:
  - File and Printer Sharing
  - Windows Management Instrumentation (WMI)
- Enable the following Windows 7 services:
  - Remote Registry
  - Windows Management Instrumentation

Create the virtual machine for the vCOps for View Adapter server

The vCOps for View Adapter server is a Windows Server 2008 R2 computer that gathers information from several sources related to View performance. The server is a required component of the vCOps for View platform.

The specifications for the server vary based on the number of desktops being monitored. Refer to the *vCenter Operations Manager for View Integration Guide* for detailed information about the resource requirements for the vCOps for View adapter server. Refer to the list of documents in Appendix C of this document for more information.

Install the vCOps for View Adapter software

Install the vCOps for View Adapter software on the server prepared in the previous step. Refer to the *vCenter Operations Manager for View Integration Guide* for detailed information about the permissions needed by the View Adapter within the components that it monitors. Refer to the list of documents in Appendix C of this document for more information.

Import the vCOps for View PAK File

The vCOps for View PAK file provides View specific dashboards for vCOps. The PAK file is located in the Program Files\VMware\vCenter Operations\View Adapter folder on the vCOps for View Adapter server, and is installed using the main vCOps web interface.

Refer to the *vCenter Operations Manager for View Integration Guide* for detailed instructions on how to install the PAK file and access the vCOps for View dashboards. Refer to the list of documents in Appendix C of this document for more information.

Verify vCOps for View functionality

Upon configuration of all required components of the vCOps for View platform, the functionality of the vCOps for View should be verified prior to deployment into production. Refer to the *vCenter Operations Manager for View Integration Guide* for detailed instructions on how to navigate the vCOps for View dashboard and observe the operation of the View environment. Refer to the list of documents in Appendix C of this document for more information.

Summary

In this chapter, we presented the requisite steps required to deploy and configure the various aspects of the VSPEX solution, which included both the physical and logical components. At this point, you should have a fully functional VSPEX solution. The following chapter covers post-installation and validation activities.

Chapter 6 Validating the Solution

This chapter presents the following topics:

Overview 186

Post-install checklist 187

Deploy and test a single virtual desktop 187

Verify the redundancy of the solution components 187

Provision remaining virtual desktops 188

Overview

This chapter provides a list of items that you should review after configuring the solution. The goal of this chapter is to verify the configuration and functionality of specific aspects of the solution and ensure that the configuration supports core availability requirements. Table 37 describes the tasks to be completed.

Table 37. Tasks for testing the installation

Task	Description	Reference
Post install checklist	Verify that adequate virtual ports exist on each vSphere host virtual switch.	vSphere Networking
	Verify that each vSphere host has access to the required datastores and VLANs.	vSphere Storage Guide vSphere Networking
	Verify that the vMotion interfaces are configured correctly on all vSphere hosts.	vSphere Networking
Deploy and test a single virtual desktop	Deploy a single virtual machine using the vSphere interface by utilizing the customization specification.	vCenter Server and Host Management vSphere Virtual Machine Management
Verify redundancy of the solution components	Restart each storage processor in turn and ensure that LUN connectivity is maintained.	Verify the redundancy of the solution components provides the steps
	Disable each of the redundant switches in turn and verify that the vSphere host, virtual machine, and storage array connectivity remains intact.	Vendor's documentation
	On a vSphere host that contains at least one virtual machine, enable maintenance mode and verify that the virtual machine can successfully migrate to an alternate host.	vCenter Server and Host Management
Provision remaining virtual desktops	Provision desktops using View Composer linked clones.	VMware Horizon View 5.3 Administration

Post-install checklist

The following configuration items are critical to functionality of the solution, and should be verified prior to deployment into production. On each vSphere server used as part of this solution, verify that:

- The vSwitches hosting the client VLANs are configured with sufficient ports to accommodate the maximum number of virtual machines it can host.
- All the required virtual machine port groups are configured and each server has access to the required VMware datastores.
- An interface is configured correctly for vMotion using the material in the *vSphere Networking* guide. Refer to the list of documents in Appendix C of this document for more information.

Deploy and test a single virtual desktop

Deploy a virtual machine to verify the operation of the solution and the procedure completes as expected. Ensure the virtual machine has been joined to the applicable domain, has access to the expected networks, and that it is possible to log in.

Verify the redundancy of the solution components

To ensure that the various components of the solution maintain availability requirements, it is important to test specific scenarios related to maintenance or hardware failure.

1. Restart each VNX Storage Processor in turn and verify that the connectivity to VMware datastores is maintained during the operation. Complete the following steps:
   a. Log into control station with administrator privileges
   b. Navigate to **/nas/sbin**
   a. Restart SPA: use command **./navicli -h spa rebootsp**
   b. During the restart cycle, check for the presence of datastores on vSphere hosts
   c. When the cycle completes, restart SPB: **./navicli -h spb rebootsp**
2. Perform a failover of each VNX Data Mover in turn and verify that connectivity to VMware datastores is maintained and that connections to CIFS file systems are reestablished. For simplicity, use the following approach for each Data Mover; use the Unisphere interface to restart.

3. From the control station $ prompt, execute the command **server_cpu *movername* -reboot**, where ***movername*** is the name of the Data Mover.
4. To verify that network redundancy features function as expected, disable each of the redundant switching infrastructures in turn. While each of the switching infrastructures is disabled, verify that all the components of the solution maintain connectivity to each other and to any existing client infrastructure as well.
5. On a vSphere host that contains at least one virtual machine, enable maintenance mode and verify that the virtual machine can successfully migrate to an alternate host.

Provision remaining virtual desktops

Complete the following steps to deploy virtual desktops using View Composer in the VMware Horizon View console:

1. Create an automated desktop pool.
2. Specify the preferred **User Assignment**:
   a. **Dedicated**: Users receive the same desktop every time they login to the pool.
   b. **Floating**: Users receive desktops picked randomly from the pool each time they log in.
3. Specify View Composer linked clones.
4. Specify a value for the Pool ID.
5. Configure Pool Settings as required.
6. Configure Provisioning Settings as required.
7. Accept default values for View Composer Disks or edit as required.
   a. If View Persona Management is used, select **Do not redirect Windows profile** in the **Persistent Disk** section, as shown in Figure 52.
   b. Configure the Active Directory Group Policy for VMware View Persona Management.

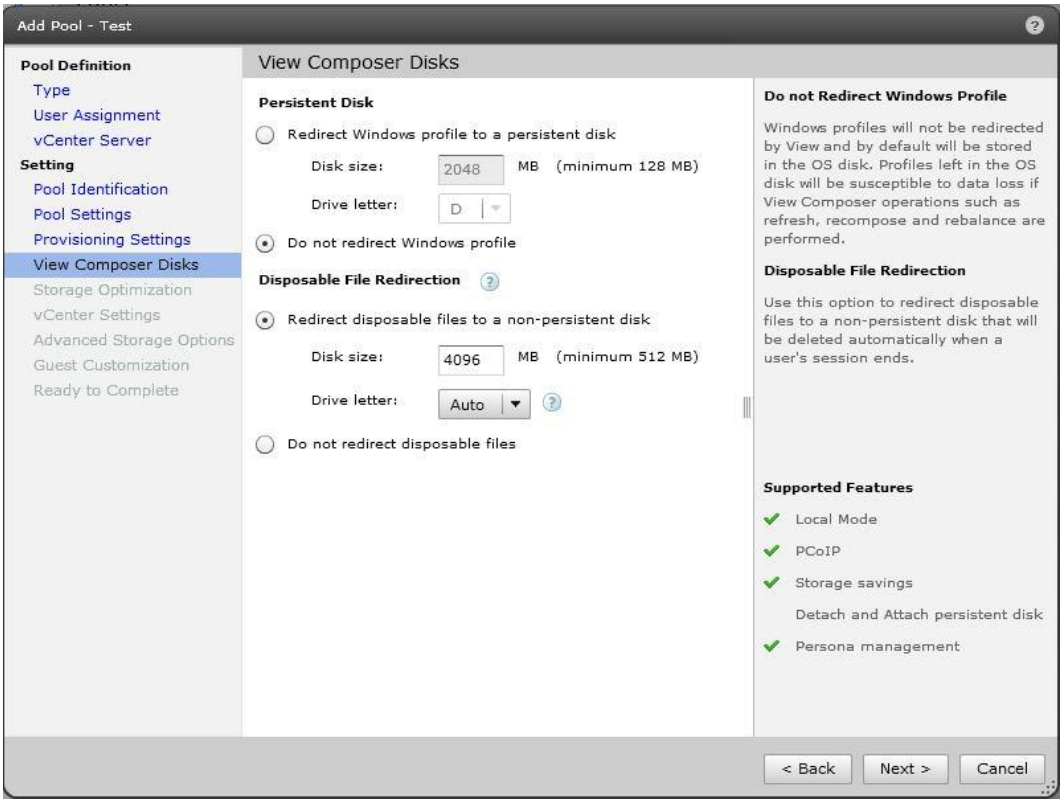
Add Pool - Test
Pool Definition
Type
User Assignment
vCenter Server
Setting
Pool Identification
Pool Settings
Provisioning Settings
View Composer Disks
Storage Optimization
vCenter Settings
Advanced Storage Options
Guest Customization
Ready to Complete
View Composer Disks
Persistent Disk
Redirect Windows profile to a persistent disk
Disk size: 2048 MB (minimum 128 MB)
Drive letter: D
Do not redirect Windows profile
Disposable File Redirection
Redirect disposable files to a non-persistent disk
Disk size: 4096 MB (minimum 512 MB)
Drive letter: Auto
Do not redirect disposable files
Do not Redirect Windows Profile
Windows profiles will not be redirected by View and by default will be stored in the OS disk. Profiles left in the OS disk will be susceptible to data loss if View Composer operations such as refresh, recompose and rebalance are performed.
Disposable File Redirection
Use this option to redirect disposable files to a non-persistent disk that will be deleted automatically when a user's session ends.
Supported Features
Local Mode
PCoIP
Storage savings
Detach and Attach persistent disk
Persona management
< Back
Next >
Cancel

Figure 52. View Composer Disks page

8. Check **Select separate datastores for replica and OS disk**.
9. Select the appropriate parent virtual machine, virtual machine snapshot, folder, vSphere hosts or clusters, vSphere resource pool, and linked clone and replica disk datastores.
10. Enable host caching for the desktop pool and specify cache regeneration blackout times.
11. Specify image customization options as required.
12. Complete the pool creation process to initiate the creation of the virtual desktop pool.

Appendix A Bills of Materials

This appendix presents the following topics:

Bill of Materials for 500 virtual desktops 192

Bill of Material for 1,000 virtual desktops 194

Bill of Material for 2,000 virtual desktops 196

Bill of Materials for 500 virtual desktops

Component		Solution for 500 virtual machines
VMware vSphere Servers	CPU	1 x vCPU per virtual machine 8 x vCPUs per physical core
		500 x vCPUs Minimum of 63 physical cores
	Memory	2 GB RAM per desktop
		2 GB RAM reservation per vSphere host
	Network – FC option	2 x 4/8 GB FC HBAs per server
	Network – 1 Gb option	6 x 1 GbE NICs per server
	Note: To implement the VMware vSphere High Availability (HA) feature and to meet the listed minimum requirements, the infrastructure should have at least one additional server beyond the number needed to meet the minimum requirements.	
Brocade Network infrastructure	Fibre Channel Block network	2 x Brocade 6510 Fibre Channel switches 4 x 4/8 Gb FC ports for VNX Backend (Two per SP) 2 x 4/8 Gb FC ports per vSphere Server for storage traffic 2 x 10 GbE ports per vSphere Server for data traffic
	10 GbE file storage network	2x Brocade VDX 6740 Ethernet Fabric switches 2 x 10GbE ports per vSphere server 2 x 10GbE ports per data mover for data
	Management network	1 x 1 GbE ports per SP controller 1 x 1 GbE port(s) per vSphere Server
	Note: When choosing the Fibre Channel option for storage, you will still need to choose one of the IP network options to have full connectivity.	
EMC Next-Generation Backup	Avamar	1 x Gen4 utility node 1 x Gen4 3.9 TB spare node 3 x Gen4 3.9 TB storage nodes

Component		Solution for 500 virtual machines
EMC VNX series storage array	Common	EMC VNX5400 2 x Data Movers (active / standby) 15 x 300 GB 15 k rpm 3.5-inch SAS drives – Core Desktops 3 x 100 GB 3.5-inch flash drives – FAST Cache 9 x 2 TB 3.5-inch NL-SAS drives (optional) – user data
	FC option	2 x 8 Gb FC ports per SP
	1 Gb network option	4 x 1 Gb I/O module for each Data Mover (each module includes four ports)
	10 Gb network option	2 x 10 Gb I/O module for each Data Mover (each module includes two ports)

Bill of Material for 1,000 virtual desktops

Component		Solution for 1,000 Virtual Machines
VMware vSphere Servers	CPU	1 x vCPU per virtual machine 8 x vCPUs per physical core
		1000 x vCPUs Minimum of 125 physical cores
	Memory	2 GB RAM per desktop
		Minimum of 2 TB RAM
	Network – FC option	2 x 4/8 GB FC HBAs per server
	Network – 1Gb option	6 x 1 GbE NICs per server
	Network – 10Gb option	3 x 10 GbE NICs per blade chassis
	Note: To implement the VMware vSphere High Availability (HA) feature and to meet the listed minimum requirements, the infrastructure should have at least one additional server beyond the number needed to meet the minimum requirements.	
Brocade Network Infrastructure	Fibre Channel Block network	2 x Brocade 6510 Fibre Channel switches 4 x 4/8 Gb FC ports for VNX Backend (Two per SP) 2 x 4/8 Gb FC ports per vSphere Server for storage traffic 2 x 10 GbE ports per vSphere Server for data traffic
	10 GbE file storage network	2x Brocade VDX 6740 Ethernet Fabric switches 2 x 10GbE ports per vSphere server 2 x 10GbE ports per data mover for data
	management network	1 x 1 GbE ports per SP controller 1 x 1 GbE port(s) per vSphere Server
	Note: When choosing the Fibre Channel option for storage, you will still need to choose one of the IP network options to have full connectivity.	

Component		Solution for 1,000 Virtual Machines
EMC Next-Generation Backup	Avamar	1 x Gen4 utility node 1 x Gen4 3.9 TB spare node 3 x Gen4 3.9 TB storage nodes
EMC VNX series storage array	Common	EMC VNX5400 2 x Data Movers (active / standby) 21 x 300 GB 15 k rpm 3.5-inch SAS drives – core desktops 3 x 100 GB 3.5-inch flash drives – FAST Cache 17 x 2 TB 3.5-inch NL-SAS drives (optional) – user data
	FC option	2 x 8 Gb FC ports per Storage Processor
	1 Gb Network option	4 x 1 Gb I/O module for each Data Mover (each module includes four ports)
	10 Gb Network option	2 x 10 Gb I/O module for each Data Mover (each module includes two ports)

Bill of Material for 2,000 virtual desktops

Component		Solution for 2,000 Virtual Machines
VMware vSphere Servers	CPU	1 x vCPU per virtual machine 8 x vCPUs per physical core
		2,000 x vCPUs Minimum of 250 Physical Cores
	Memory	2 GB RAM per desktop
		Minimum of 4 TB RAM
	Network – FC option	2 x 4/8 GB FC HBAs per server
	Network – 1 Gb option	6 x 1 GbE NICs per server
	Network – 10 Gb option	3 x 10 GbE NICs per blade chassis
	Note: To implement VMware vSphere High Availability (HA) functionality and to meet the listed minimum requirements, the infrastructure should have at least one additional server beyond the number needed to meet the minimum requirements.	
Brocade Network Infrastructure	Fibre Channel Block network	2 x Brocade 6510 Fibre Channel switches 4 x 4/8 Gb FC ports for VNX Backend (Two per SP) 2 x 4/8 Gb FC ports per vSphere Server for storage traffic 2 x 10 GbE ports per vSphere Server for data traffic
	10 GbE file storage network	2x Brocade VDX 6740 Ethernet Fabric switches 2 x 10GbE ports per vSphere server 2 x 10GbE ports per data mover for data
	management network	1 x 1 GbE ports per SP controller 1 x 1 GbE port(s) per vSphere Server
	Note: When choosing the Fibre Channel option for storage, you will still need to choose one of the IP network options to have full connectivity.	

Component		Solution for 2,000 Virtual Machines
EMC Next-Generation Backup	Avamar	See Design and Implementation Guide: EMC Backup and Recovery Options for VSPEX End User Computing for VMware Horizon View available on EMC Online Support.
EMC VNX series storage array	Common	EMC VNX5600 3 x Data Movers (active/standby) 36 x 300 GB 15 k rpm 3.5-inch SAS drives – core desktops 5 x 100 GB 3.5-inch flash drives – FAST Cache 34 x 2 TB 3.5-inch NL-SAS drives (optional) – user data
	FC option	2 x 8 Gb FC ports per Storage Processor
	1 Gb Network option	4 x 1 Gb I/O module for each Data Mover (each module includes four ports)
	10 Gb Network option	2 x 10 Gb I/O module for each Data Mover (each module includes two ports)

Appendix B Customer Configuration Data Sheet

This appendix presents the following topic:

Overview of customer configuration data sheets 200

Overview of customer configuration data sheets

Before you start the configuration, gather customer-specific network and host configuration information. The following tables provide information on assembling the required network and host address, numbering, and naming information. This worksheet can also be used as a "leave behind" document for future reference.

The *VNX File and Unified Worksheet* should be cross-referenced to confirm customer information.

Table 38. Common Server information

Server Name	Purpose	Primary IP
	Domain Controller	
	DNS Primary	
	DNS Secondary	
	DHCP	
	NTP	
	SMTP	
	SNMP	
	VMware vCenter Console	
	VMware View Connection Servers	
	Microsoft SQL Server	
	VMware vShield Manager	
	Antivirus solution management server	
	vCOps for View Adapter server	

Table 39. vSphere Server information

Server Name	Purpose	Primary IP	Private Net (storage) addresses	VMkernel IP	vMotion IP
	vSphere Host 1				
	vSphere Host 2				
	…				

Table 40. Array information

Array name	
Admin account	
Management IP	
Storage pool name	
Datastore name	
NFS Server IP	

Table 41. Brocade Network infrastructure information

Name	Purpose	IP	Subnet Mask	Default Gateway
	Ethernet switch 1			
	Ethernet switch 2			
	…			

Table 42. VLAN information

Name	Network Purpose	VLAN ID	Allowed Subnets
	Virtual Machine Networking		
	vSphere Management		
	NFS storage network		
	vMotion		

Table 43. Service accounts

Account	Purpose	Password (optional, secure appropriately)
	Windows Server administrator	
root	vSphere root	
root	Array root	
	Array administrator	
	VMware vCenter administrator	
	VMware Horizon View administrator	
	SQL Server administrator	
	VMware vCOps administrator	
	VMware vShield Manager administrator	

Appendix C References

This appendix presents the following topic:

References 204

References

EMC documentation

The following documents, located on EMC Online Support provide additional and relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your EMC representative.

EMC Avamar 7.0 Administrator Guide

EMC Avamar 7.0 Operational Best Practices

EMC VNX5400 Unified Installation Guide

EMC VNX5600 Unified Installation Guide

EMC VSI for VMware vSphere: Storage Viewer Product Guide

EMC VSI for VMware vSphere: Unified Storage Management Product Guide

Deployment and Configuration Guide: vCenter Operations Manager 5

vCenter Operations Manager for View Integration Guide

VNX Block Configuration Worksheet

VNX File and Unified Worksheet

- *VNX FAST Cache:* A Detailed Review

Unisphere System Getting Started Guide

Using EMC VNX Storage with VMware vSphere TechBook

Deploying Microsoft Windows 7 Virtual Desktops with VMware View : Applied Best Practices white paper

PowerPath/VE for VMware vSphere Installation and Administration Guide

PowerPath Viewer Installation and Administration Guide

EMC Technical Note: Avamar Client for Windows on Virtual Desktops (P/N 300-011-893)

EMC VNX Unified Best Practices for Performance: Applied Best Practices Guide

Sizing EMC VNX Series for VDI Workload

EMC Infrastructure for VMware View 5.1: EMC VNX Series (FC), VMware vSphere 5.0, VMware View 5.1, VMware View Storage Accelerator, VMware View Persona Management, and VMware View Composer 3.0 Reference Architecture

EMC Infrastructure for VMware View 5.1 — EMC VNX Series (FC), VMware vSphere 5.0, VMware View 5.1, VMware View Storage Accelerator,

VMware View Persona Management, and VMware View Composer 3.0 Proven Solutions Guide

EMC Infrastructure for VMware View 5.1 — EMC VNX Series (NFS), VMware vSphere 5.0, VMware View 5.1, VMware View Storage Accelerator, VMware View Persona Management, and VMware View Composer 3.0 Reference Architecture

EMC Infrastructure for VMware View 5.1 — EMC VNX Series (NFS), VMware vSphere 5.0, VMware View 5.1, VMware View Storage Accelerator, VMware View Persona Management, and VMware View Composer 3.0 Proven Solutions Guide

EMC Infrastructure for Citrix XenDesktop 5.6, EMC VNX Series (NFS), VMware vSphere 5.0, Citrix XenDesktop 5.6, and Citrix Profile Manager 4.1 Reference Architecture

EMC Infrastructure for Citrix XenDesktop 5.6 — EMC VNX Series (NFS), VMware vSphere 5.0, Citrix XenDesktop 5.6, and Citrix Profile Manager 4.1 Proven Solutions Guide

EMC Infrastructure for Citrix XenDesktop 5.5 (PVS) — EMC VNX Series (NFS), Citrix XenDesktop 5.5 (PVS), XenApp 6.5, and XenServer 6 Reference Architecture

EMC Infrastructure for Citrix XenDesktop 5.5 (PVS) EMC VNX Series (NFS), Citrix XenDesktop 5.5 (PVS), XenApp 6.5, and XenServer 6 Proven Solution Guide

EMC Infrastructure for Citrix XenDesktop 5.5: EMC VNX Series (NFS), Citrix XenDesktop 5.5, XenApp 6.5, and XenServer 6

EMC Infrastructure for Citrix XenDesktop 5.5 — EMC VNX Series (NFS), Citrix XenDesktop 5.5, XenApp 6.5, and XenServer 6 Proven Solution Guide

Brocade Documentation

Brocade VDX Switches and VCS Fabrics related documentation can be found as stated below:

Brocade VDX 6740/6740T/ 6740T-1G Switch Data sheet

http://www.brocade.com/downloads/documents/data_sheets/product_data_sheets/vdx-6740-ds.pdf

Hardware Reference Manual

Brocade VDX 6740 Hardware Reference Manual
http://www.brocade.com/downloads/documents/product_manuals/B_VDX/VDX6740_VDX6740T_HardwareManual.pdf

-

Brocade Network OS (NOS) Guides

Network OS Administrator's Guide Supporting Network OS v4.1.0
http://www.brocade.com/downloads/documents/product_manuals/B_VDX/NOS_AdminGuide_v410.pdf

Network OS Command Reference Supporting Network OS v4.1.0
http://www.brocade.com/downloads/documents/product_manuals/B_VDX/NOS_CommandRef_v301.pdf

http://www.brocade.com/downloads/documents/product_manuals/B_VDX/NOS_CommandRef_v410.pdf

Brocade Network OS (NOS) Software Licensing Guide v4.1.0
http://www.brocade.com/downloads/documents/product_manuals/B_VDX/NOS_LicensingGuide_v410.pdf

The Brocade Network Operating System (NOS) Release notes can be found at http://my.brocade.com

Brocade 65xx Switches and FOS Fabrics related documentation can be found as stated below:

Product Data Sheets for Brocade VDX 6510 Series of switches:

http://www.brocade.com/downloads/documents/data_sheets/product_data_sheets/6510-switch-ds.pdf

Hardware Reference Manual

http://www.brocade.com/downloads/documents/product_manuals/B_SAN/B6510_HardwareManual.pdf

Brocade Fabric OS (FOS) Guides

Fabric OS Administrator's Guide Supporting Network OS v7.2.0

http://www.brocade.com/downloads/documents/product_manuals/B_SAN/FOS_AdminGd_v720.pdf

Fabric OS Command Reference Supporting Network OS v7.2.0

http://www.brocade.com/downloads/documents/product_manuals/B_SAN/FOS_CmdRef_v720.pdf

Brocade 6510 QuickStart Guide
http://www.brocade.com/downloads/documents/product_manuals/B_SAN/B6510_QuickStartGuide.pdf

SAN Fabric Administration Best Practices Guide
http://www.brocade.com/downloads/documents/best_practice_guides/san-admin-best-practices-bp.pdf

The Brocade Fabric Operating System (FOS) Release notes can be found at http://my.brocade.com

Other documentation

The following documents, located on the VMware website, provide additional and relevant information:

Installing and Administering VMware vSphere Update Manager

Preparing vCenter Server Databases

Preparing the Update Manager Database

vCenter Server and Host Management

View 5.3 Administration Guide

View 5.3 Architecture and Planning Guide

View 5.3 Installation Guide

View 5.3 Integration Guide

View 5.3 Profile Migration Guide

View 5.3 Security Guide

View 5.3 Upgrades Guide

VMware vCenter Operations Manager for View Integration Guide

VMware vCenter Operations Manager Administration Guide

VMware vCenter Operations Manager Installation Guide

VMware View Optimization Guide for Windows 7

vShield Administration Guide

vShield Quick Start Guide

vSphere Resource Management

vSphere Storage APIs for Array Integration (VAAI) Plug-in

vSphere Installation and Setup Guide

vSphere Networking

vSphere Storage Guide

vSphere Virtual Machine Administration

vSphere Virtual Machine Management

For documentation on Microsoft SQL Server, refer to the following Microsoft websites:

www.microsoft.com

technet.microsoft.com

msdn.microsoft.com

Appendix D About VSPEX

This appendix presents the following topic:

About VSPEX 210

About VSPEX

EMC has joined forces with the industry's leading providers of IT infrastructure to create a complete virtualization solution that accelerates deployment of cloud infrastructure. Built with best-in-class technologies, VSPEX enables faster deployment, more simplicity, greater choice, higher efficiency, and lower risk. Validation by EMC ensures predictable performance and enables customers to select technology that uses their existing IT infrastructure while eliminating planning, sizing, and configuration burdens. VSPEX provides a proven infrastructure for customers looking to gain simplicity that is characteristic of truly converged infrastructures while at the same time gaining more choice in individual stack components.

VSPEX solutions are proven by EMC and packaged and sold exclusively by EMC channel partners. VSPEX provides channel partners more opportunity, faster sales cycle, and end-to-end enablement. By working even more closely together EMC and its channel partners can now deliver infrastructure that accelerates the journey to the cloud for even more customers.